

Outlook

PROPOSTA FINAL | PE 01/2025 | PERFIL COMP

De Vinicius Araldi <vinicius.araldi@perfil.inf.br>
 Data Sex, 28/03/2025 16:08
 Para Licitações <licitacoes@paranaprojetos.org.br>

Cc Francisco A. Jarschel <francisco@perfil.inf.br>

4 anexos (24 MB)

0. Proposta Perfil.pdf; Declarações do fabricante.zip; Documentos Tecnicos.zip; Documentos de habilitação.zip;

Prezados, boa tarde!

Segue em anexo proposta, declarações, documentos técnicos e documentos de habilitação referente ao pregão eletrônico 01/2025.

Fico a disposição para eventuais dúvidas.

Atenciosamente,

Vinicius Araldi | Auxiliar Comercial | 54.2628-8305



https://outlook.office.com/mail/inbox/id/AAMkAGQxNmMyYTMwLTNhNTgtNDJhNS1hMDQwLTIzY2Y4MmNjZTBmMgBGAAAAAACK8UTNwzZYS... 1/1



D&LLTechnologies

Declaração do Fabricante

A **DELL COMPUTADORES DO BRASIL LTDA.** ("Dell"), inscrita no CNPJ sob o n. 72.381.189/0001-10, na qualidade de fabricante do(s) equipamento(s) de marca Dell (abaixo identificado(s)), ofertado(s) pela empresa Perfil Computacional, no certame licitatório n. PE 01/2025 promovido pelo PARANA PROJETOS, vem, através desta, declarar que:

- o(s) modelo(s) Dell PowerEdge R760xs possui(em) garantia de 60 meses, on-site, com atendimento telefônico 24 horas por dia, 7 dias na semana.

Declaramos, ainda, que:

- Os equipamentos por nós fabricados serão novos, sem uso e não são produtos descontinuados.

- A Perfil Computacional está autorizada a comercializar os equipamentos propostos para esse certame.

Eldorado do Sul, 28 de março de 2025



Dell Computadores do Brasil Ltda Luciano Valadares – Executivo de Contas

DELL Computadores do Brasil Ltda. Av. Industrial Belgraf, 400 . Eldorado do Sul / RS . Geral : 51 3481 5500 Fax : 51 3481 5458



D&LLTechnologies

Eldorado do Sul, 27 de março de 2025

À Perfil Computacional A/C Sr. Igor S. Reolon

Ref.: PARANÁ PROJETOS – Pregão Eletrônico 01/2025

DECLARAÇÃO TÉCNICA

DELL COMPUTADORES DO BRASIL LTDA. ("Dell"), inscrita no CNPJ/MF sob o n° 72.381.189/0001-10, com sede na Av. Industrial Belgraf, 400 – Medianeira – CEP 92990-000, Eldorado do Sul/RS, com o objetivo de complementar as informações que não constam no Catálogo Técnico Oficial do(s) produto(s) abaixo ofertado(s), vem, através da presente, declarar o que segue:

Objeto: Dell PowerEdge R760xs

- A BIOS é desenvolvida pela Dell Computadores do Brasil, detentora dos direitos copyright.

- Placa Mãe da fabricante DELL sendo projetada e desenvolvida especificamente para os modelos ofertado, não sendo uma placa de livre comercialização no mercado.

- Os manuais, drivers, firmwares e atualizações dos produtos Dell são disponibilizados no site do fabricante: www.dell.com.br/suporte.

- Todos os equipamentos e seus opcionais são integrados em fábrica, sem quaisquer adaptações.

- Os equipamentos ofertados são novos e estão em atual linha de produção.



Dell Computadores do Brasil Ltda Luciano Valadares – Executivo de Contas

DELL Computadores do Brasil Ltda. Av. Industrial Belgraf , 400 . Eldorado do Sul / RS . Geral : 51 3481 5500 Fax : 51 3481 5458

Perfilcomp



ANEXO III - TERMO DE DECLARAÇÃO

Ao SERVIÇO SOCIAL AUTÔNOMO PARANÁ PROJETOS

Rua Inácio Lustosa, 700 - Bloco A - Térreo - Curitiba - Paraná.

Ref.: PREGÃO ELETRÔNICO PARANÁ PROJETOS N.º 01/2025 (PROTOCOLO: 22.951.206-4) - AQUISIÇÃO DE SERVIDOR, COM PLANO DE GARANTIA ESTENDIDA DO FABRICANTE, EM CONFORMIDADE COM AS ESPECIFICAÇÕES TÉCNICAS MÍNIMAS EXIGIDAS, PARA ATENDER AS DEMANDAS DO PARANÁ PROJETOS.

A empresa Perfil Computacional Ltda, CNPJ 02.543.216/0011-09, com sede no endereço Av. Acesso Rodoviário, SN, Quadra 06 LM 01 Quadra 01 L-M18 AM23 Sala 107, Terminal Intermodal da Serra, Serra – ES, CEP: 29161-376, Telefone (54) 2628-8300, e-mail perfil@perfil.inf.br, por seu representante abaixo assinado, DECLARA QUE:

I) Examinou cuidadosamente o edital, inteirou-se de todos os seus detalhes e com eles concorda, aceita todos os seus termos e condições e a eles desde já se submete;

II) Todas as dúvidas ou questionamentos formulados foram devidamente esclarecidos, bem como recebeu todos os elementos e informações para cumprimento das obrigações objeto da licitação;

III) Nos valores constantes da proposta estão incluídas todas as despesas decorrentes da execução do contrato, tais como materiais, mão de obra, custos diretos e indiretos, despesas com encargos sociais, previdenciários, trabalhistas, seguros, taxas, tributos e contribuições de qualquer natureza ou espécie, transportes e quaisquer outros encargos necessários à perfeita execução do objeto do contrato;

IV) Para os fins de participação nesta licitação, declara fundamentalmente que:

a) está ciente, conhece e entende os termos das leis anticorrupção brasileiras ou de quaisquer outras aplicáveis sobre o objeto ora licitado;

b) não foi condenada pelas práticas previstas na Lei 12.846/2013;

c) seus sócios, diretores, administradores, empregados, assessores, prepostos e colaboradores não cometerão, auxiliarão, incitarão ou instigarão terceiros a cometerem atos ilícitos, que incluem oferecer, conceder, requerer ou aceitar pagamentos, doações, compensação, benefícios ou quaisquer outras vantagens indevidas e/ou ilegais para si ou para terceiros, bem como o desvio de finalidade do presente contrato, ou atos lesivos expressamente previstos na Lei Federal n.º 12.846/13, que constituam prática ilegal ou de corrupção, seja de forma direta ou indireta quanto ao objeto deste contrato;

V) A signatária não se encontra suspensa de licitar ou contratar com o PARANÁ PROJETOS e inexiste empregado, diretor e conselheiro do PARANÁ PROJETOS na composição societária da empresa.

VI) Conhece e não apresenta as restrições contidas no Cadastro de Empregadores que tenham submetido trabalhadores as condições análogas à de escravo - "Lista Suja" - nos termos da Portaria MTE nº 4/2016, bem como não emprega menor de dezoito anos em trabalho noturno, perigoso ou insalubre e não emprega menor de dezesseis anos, salvo na condição de aprendiz, em conformidade ao art. 7º, XXXIII, da Constituição Federal.

VII) Está ciente que a licitação NÃO SERÁ HOMOLOGADA, caso vencedora do certame, seja constatada sua inclusão no referido cadastro, por meio de consulta.

Serra, 28 de março de 2025. IGOR SIDNEI REOLON:8051

Assinado de forma digital por IGOR SIDNEI REOLON:80512771049 Dados: 2025.03.28 10:32:24 -03'00'

Igor Sidnei Reolon Representante Legal – CEO CPF nº 805.127.710-49 | RG nº 1067955946 Perfil Computacional Ltda CNPJ: 02.543.216/0011-09

0800 721 0675

📓 Nossas unidades: Farroupilha RS | Florianópolis SC | Curitiba PR | Barueri SP | Brasília DF | Serra ES

2771049







ANEXO VI - FORMULÁRIO DE CREDENCIAMENTO - LICITAÇÃO

Declaro, sob as penas da lei, que as informações abaixo são verdadeiras, pelos quais firmo a presente.

FORMULARIO DE CREDENCIAMENTO - LICITAÇÃO

IDENTIFICAÇÃO RAZÃO SOCIAL: Perfil Computacional Ltda NOME FANTASIA: Perfilcomp NOME DO REPRESENTANTE PARTICIPANTE DA LICITAÇÃO: Igor Sidnei Reolon E-MAIL: perfil@perfil.inf.br NOME DO RESPONSAVEL PELA ASSINATURA DO INSTRUMENTO: Rodrigo Alves Soares **CARGO:** Gerente Comercial – Procurador **CNPJ:** 02.543.216/0011-09 **INSCRIÇÃO ESTADUAL:** 083.644.61-0

PORTE DA EMPRESA:

□ MEI (MICROEMPREENDEDOR INDIVIDUAL) □ ME (MICROEMPRESA) □ EPP (EMPRESA DE PEQUENO PORTE) X OUTROS PORTES OBS.:

ENQUADRAMENTO TRIBUTÁRIO:

□ SIMPLES NACIONAL X LUCRO REAL □ LUCRO PRESUMIDO □ LUCRO ARBITRADO □ SEM FINS LUCRATIVOS □ OUTROS

ENDEREÇO DA EMPRESA

ENDEREÇO: Av. Acesso Rodoviário NÚMERO: SN COMPLEMENTO: Quadra 06 LM 01 Quadra 01 L-M18 AM23 Sala 107 BAIRRO: Terminal Intermodal da Serra CEP: 29161-376 CIDADE: Serra ESTADO: ES

ENDEREÇO ELETRÔNICO (E-mail): perfil@perfil.inf.br

ENDEREÇO PARA CORRESPONDÊNCIA

ENDEREÇO: Rua Barão do Rio Branco NÚMERO: 459 COMPLEMENTO: Sala 20 BAIRRO: Centro **CEP:** 95170-404 **CIDADE:** Farroupilha ESTADO: RS

E-MAIL PESSOAL DO RESPONSÁVEL PELA ASSINATURA DO CONTRATO: rodrigo@perfil.inf.br TELEFONE: (54) 2628-8308

Serra, 28 de março de 2025.

IGOR SIDNEL REOLON:80512 REOLON:80512771049

771049

Assinado de forma digital por IGOR SIDNEI Dados: 2025.03.28 10:32:06 -03'00'

Igor Sidnei Reolon Representante Legal – CEO CPF nº 805.127.710-49 | RG nº 1067955946 Perfil Computacional Ltda CNPJ: 02.543.216/0011-09

0800 721 0675

📱 Nossas unidades: Farroupilha RS | Florianópolis SC | Curitiba PR | Barueri SP | Brasília DF | Serra ES







JUSTIÇA ELEITORAL TRIBUNAL REGIONAL ELEITORAL DO RIO GRANDE DO SUL

Comunicações - Serviços - 0012050-28.2020.6.21.8000 Atestado - doc. SEI n. 0348210.

ATESTADO DE CAPACIDADE TÉCNICA

Nº 20/2020

Atestamos, para os devidos fins, que a empresa Perfil Computacional Ltda., CNPJ n. 02.543.216/0001-29, situada na Rua Barão do Rio Branco, n. 459, sala 20, Bairro Centro em Farroupilha-RS, CEP 95.180-000, forneceu a este Tribunal, CNPJ n. 05.885.797/0001-75:

1) Solução de armazenamento de dados baseado em SDS (Software Defined Storage) com capacidade mínima de armazenamento de 40TB líquidos, com replicação de 3 (três) cópias, com performance de I/O de 20.000 IOPS de escrita e 50.000 IOPS de leitura, mediante fornecimento de hardware, podendo aproveitar 4(quatro) servidores Dell PowerEdge R640 do TRE, com instalação e configuração, treinamento e subscrição de software SDS (CEPH) para 36 meses de serviço. Para tanto, forneceu os seguintes itens:

- 04 (quatro) servidores Dell Poweredge R740XD Server, BCC (210-ALCR), com as seguintes características: chassis with up to 24 x 2.5" hard drives (321-BCPY), 02 (dois) intel xeon gold 5120 de 2,2GHZ, 14 núcleos/28 segmentos, 10,4 GT/S, cache de 19,25MB, turbo, HT (105W), DDR4-2400 (338-BLUX), 12 (doze) 16GB RDIMM, 2666MT/S, dual rank, BCC; HBA330 controller, 12GBPS adapter, low profile (405-AANK), marca Dell EMC, fabricantre Dell EMC, modelo/versão Poweredge R740XD, conforme nota de empenho n. 2018NE001418;

- Licença de software para solução de armazenamento de dados baseado em SDS (software defined storage), por 03 anos, com capacidade mínima de armazenamento de 40TB líquidos, com replicação de 3 (três) cópias, com performance de I/O de 20.000 IOPS de escrita e 50.000 IOPS de leitura, marca Suse, fabricante Suse, modelo/versão **Suse Enterprise Storage**, conforme nota fiscal 8581.

- Prestação de serviços de instalação e configuração da solução de armazenamento (item 3.2.4 do TR). O cluster deverá ser instalado nos racks padrão 19" já disponíveis no datacenter do contratante. Devendo, ainda, ser ministrado treinamento (item 3.2.5 do TR) observando que seja do tipo hands on, por profissional certificado pelo distribuidor do SDS instalado, com no mínimo de 08 horas/aula, abrangendo manuseio, configuração e operação da solução, para turma de até 06 servidores do TRE-RS, conforme nota de empenho n. 2018NE001420.

2) Subscrição de suporte do software Proxmox Virtual Environment por 36 meses para 4 servidores biprocessados (total de 8 sockets).

- 08 (oito) subscrições de suporte do software Proxmox Virtual Environment, pelo período de 36 meses, com os

seguintes requisitos: permitir a abertura de chamados via interface web, permitir a abertura de até 10 chamados por ano, permitir acesso ao repositório de software enterprise (enterprise repository), possuir tempo de solução de 1 dia útil após abertura do chamado, o serviço de suporte deverá ser prestado pela distribuidora do software, marca Proxmox, fabricante Proxmox, conforme nota de empenho n. 2018NE001419.

Demais especificações conforme Edital do Pregão n. 60/2018, Processo Eletrônico Administrativo n. 144/2018 e Processo SEI n. 0003523-24.2019.6.21.8000.

A empresa demonstrou boa capacidade técnica não constando em nossos registros fato que a desabone.

Porto Alegre, 15 de julho de 2020.

JOSÉ ATILIO BENITES LOPES COORDENADORIA DE MATERIAL E PATRIMÔNIO



Documento assinado eletronicamente por **Jose Atilio Benites Lopes**, **Coordenador**, em 15/07/2020, às 19:23, conforme art. 1°, § 2°, III, "b", da Lei 11.419/2006.



A autenticidade do documento pode ser conferida no site https://sei.tre-rs.jus.br/sei/controlador_externo.php? acao=documento_conferir&id_orgao_acesso_externo=0 informando o código verificador **0348210** e o código CRC **216E08B6**.

Rua Duque de Caxias, 350 - Bairro Centro - Porto Alegre/RS - CEP 90010-280 www.tre-rs.jus.br - Fone: (51) 3294 8314



ESTADO DE SANTA CATARINA SECRETARIA DE ESTADO DA FAZENDA DIRETORIA DE ADMINISTRAÇÃO TRIBUTÁRIA GERÊNCIA DE SISTEMAS DE ADMINISTRAÇÃO TRIBUTÁRIA



ATESTADO DE CAPACIDADE TÉCNICA

A **SECRETARIA DE ESTADO DA FAZENDA DO ESTADO DE SANTA CATARINA**, inscrita no CNPJ sob o nº 82.951.310/0001-56, atesta para os devidos fins, que a empresa **Perfil Computacional Ltda**, com sua filial na Av. Rio Branco, nº 404, Torre II, sl 908-A, Centro, Florianópolis, SC, CEP 88015-200, inscrita no CNPJ Nº: 02.543.216/0006-33 - Insc.Estadual: 257927212, participante do PREGÃO ELETRÔNICO Nº 0055/2024 da Secretaria de Estado da Fazenda de Santa Catarina, promovido pela SECRETARIA DE ESTADO DA ADMINISTRAÇÃO, entregou os equipamentos/materiais/solução abaixo discriminados:

Maiores informações podem ser encontradas nos documentos relacionados ao: PREGÃO ELETRÔNICO Nº 0055/2024 CONTRATO SEF 027/2024

ltem	Serviço	Qtd	Métrica
7	Servidor para Master Node (Dell		
	PowerEdge R660)		
		3	Unidade
8	Servidor para Worker Node (Dell		
	PowerEdge R660)		
		8	Unidade
9	Servidor para Worker Node com		
	GPU (Dell PowerEdge R760 com		
	2x GPU's NVIDIA L40s)	1	Unidade
10	Solução de armazenamento All-		
	Flash NVMe (5 unidades Dell		
	PowerScale F600 - 8x discos de	1	Unidade
	15.36TB cada unidade)		
11	Switch Tor (Dell PowerSwitch		
	S5232F)		
		2	Unidade
12	Serviço de instalação,	1	Unidade
	configuração e tuning dos		
	equipamentos		
1	1	1	1

Gerência de Sistemas de Administração Tributária

Rua Tenente Silveira, 60 - 3º andar - Centro – Florianópolis - SC CEP 88.010-300 CNPJ: 82 951 310 0001-56 | Fone: (48) 3664-5711 | <u>www.sef.sc.gov.br</u>



ESTADO DE SANTA CATARINA SECRETARIA DE ESTADO DA FAZENDA DIRETORIA DE ADMINISTRAÇÃO TRIBUTÁRIA GERÊNCIA DE SISTEMAS DE ADMINISTRAÇÃO TRIBUTÁRIA



Atestamos ainda, que a empresa cumpriu e vem cumprindo satisfatoriamente com as obrigações contratuais e demais condições estabelecidas, durante o prazo de garantia dos equipamentos, nada constando que a desabone até a presente data.

Por ser verdade, firmo o presente.

Florianópolis, data da assinatura digital.

OMAR ROBERTO AFIF ALEMSAN:31889654949 Assinado de forma digital por OMAR ROBERTO AFIF ALEMSAN:31889654949 Dados: 2024.11.29 17:09:34 -03'00'

Omar Roberto Afif Alemsan Gerente de Sistemas de Administração Tributária

Gerência de Sistemas de Administração Tributária Rua Tenente Silveira, 60 - 3º andar - Centro – Florianópolis - SC CEP 88.010-300 CNPJ: 82 951 310 0001-56 | Fone: (48) 3664-5711 | <u>www.sef.sc.gov.br</u>





Prefeitura do Município de Foz do Iguaçu

ESTADO DO PARANÁ

ATESTADO DE CAPACIDADE TÉCNICA

Foz do Iguaçu, 15 de setembro de 2023.

Atestamos para os devidos fins que a Empresa PERFIL COMPUTACIONAL LTDA, com sede na Av. Acesso Rodoviário, SN, Terminal, Quadra 06 LM 01 Quadra 01 L-M18 AM23 SALA 107, Intermodal da Serra, Serra – ES, inscrita no CNPJ sob o n° 02.543.216/0011-09, forneceu ao Município de Foz do Iguaçu - PR, inscrito sob o CNPJ n°76.206.606/0001-40,os itens que segue abaixo, para o Pregão Eletrônico 213/2022.

3 x Servidor de Virtualização de 2 Soquetes Dell EMC PowerEdge R6525

1 x Solução de Armazenamento (storage) All Flash Dell EMC Unity XT 480F com 480 TiB

2 x Switch de Produção Dell EMC PowerSwitch S5248F

1 x Switch de Gerenciamento Dell EMC PowerSwitch S3148T

1 x Appliance para Solução de Backup Dell EMC PowerProtect DD6900

1 x Licenciamento Dell EMC BAse Data Domain , DD Bust e DD Replication

2x Nobreaks – Sistema UPS de no mínimo 10KW

1 x Rack 19" para acondicionar os hardwares APC Netshelter AR3100

1 x VMware Cloud Foundation 4 Standard for External Storage Per CPU (CF4-STD-ES-C) e vCenter Server Standard

3x Microsoft Windows Server Data Center 2022

2500 x Microsoft Windows Server CAL de acesso, por dispositivo

100 x Microsoft Windows Server CAL RDS, por usuário

12 x Microsoft SQL Server Standard 2022 com Assurance 36 meses

1 x Software Solução de Backup Dell EMC Data ProtectionSuite

Instalação e configuração Vmware(vCenter ,vSphere ,vRealizeOperations ,vRealizeLifecycle Manager ,vRealizeLogInsight ,NSX-T(Pro))

Migração de Maquinas Virtuais VMWare para novo Cluster

Atualização Microsoft Windows DomaisnControllers

Este documento foi assinado eletronicamente por vários signatários.

Para verificar as assinaturas vá ao site https://sistemas.pmfi.pr.gov.br/rp/sidpublico/verificar e utilize o código 801507b5-6a67-444b-9457-828315e3f032.

Inserido ao protocolo **22.951.206-4** por: **Pedro Henrique Golin Linhares** em: 31/03/2025 12:49. A autenticidade deste documento pode ser validada no endereço: https://www.eprotocolo.pr.gov.br/spiweb/validarDocumento com o código: 1d3e5a42275b5642ba40f346676ed94.



Prefeitura do Município de Foz do Iguaçu

ESTADO DO PARANÁ

ATESTADO DE CAPACIDADE TÉCNICA

Foz do Iguaçu, 15 de setembro de 2023.

Migração Microsoft Windows Server para novos servidores virtuais

Configuração ambiente físico Microsoft Cluster SQL Server

Instalação física e lógica de todos os equipamentos fornecidos

Instalação, configuração e criação de rotinas de backup de ferramenta de Backup Dell EMC Networker

Os produtos foram entregues dentro do prazo e condições, atendendo aos requisitos técnicos exigidos, inclusive em relação a assistência técnica, durante o prazo de garantia, nada havendo a desabonar a Capacidade Técnica do referido fornecedor.

Por expressão de verdade, assina o presente atestado.

O(s) nome(s) indicado(s) para assinatura: Renato Vieira Gomes - **Diretor de Infraestrutura e Segurança da Informação - Portaria nº 75.212/2022** Evandro Ferreira - **Secretário Municipal de Tecnologia da Informação Portaria 75.659/2023**

Ao Ilmo. Sr.: **André Bellaver - Gerente de Contas** Perfil Computacional LTDA. Av. Cândido de Abreu, 776 - Centro Cívico <u>CURITIBA-PR</u> CEP 80530-000

Este documento foi assinado eletronicamente por vários signatário

Para verificar as assinaturas vá ao site https://sistemas.pmfi.pr.gov.br/rp/sidpublico/verificar e utilize o código 801507b5-6a67-444b-9457-828315e3f032.

PROTOCOLO DE ASSINATURA(S)

Tipo: ATESTADO DE CAPACIDADE TÉCNICA

Número: 5/2023

Assunto: ATESTADO DE CAPACIDADE TÉCNICA - PERFIL COMUTACIONAL LTDA

O documento acima foi proposto para assinatura eletrônica na plataforma **SID** de assinaturas. Para verificar as assinaturas clique no link:

https://sistemas.pmfi.pr.gov.br/rp/sidpublico/verificar?codigo=801507b5-6a67-444b-9457-828315e3f032&cpf=92539939953 e utilize o código abaixo para verificar se este documento é válido.

Código para verificação: 801507b5-6a67-444b-9457-828315e3f032

Hash do Documento

D6413195FDFFE40FA7A936C248CC12663580F0B860255FFA2FE6BB4CB7661E45

O(s) nome(s) indicado(s) para assinatura, bem como seu(s) status em 18/09/2023 é(são) :

RENATO VIEIRA GOMES (Signatário) - CPF: ***59782997** em 15/09/2023 8:41:56 - OK Tipo: Assinatura Eletrônica

EVANDRO FERREIRA (Signatário) - CPF: ***39939953** em 18/09/2023 7:53:16 - OK Tipo: Assinatura Eletrônica



A ASSINATURA ELETRÔNICA DESTE DOCUMENTO ESTÁ AMPARADA PELO:

DECRETO Nº 28.900, DE 20 DE JANEIRO DE 2021.

LEI Nº 4536 , DE 4 DE SETEMBRO DE 2017.

Autoriza a utilização do meio eletrônico para a gestão dos processos administrativos e de documentos de arquivo, produzidos nos termos das Leis nºs 3.971, de 17 de abril de 2012 e 4.057, de 19 de dezembro de 2012, no âmbito dos órgãos da Administração Pública Direta, Autárquica e Fundacional do Município de Foz do Iguaçu.



TERMOS DE ABERTURA E ENCERRAMENTO

Entidade:

PERFIL COMPUTACIONAL LTDA

Período da Escrituração:01/01/2023 a 31/12/2023Número de Ordem do Livro:26

CNPJ: 02.543.216/0001-29

247

TERMO DE ABERTURA			
Nome Empresarial	PERFIL COMPUTACIONAL LTDA		
NIRE	43203789542		
CNPJ	02.543.216/0001-29		
Número de Ordem	26		
Natureza do Livro	LIVRO DIARIO		
Município	FARROUPILHA		
Data do arquivamento dos atos constitutivos	19/05/1998		
Data de arquivamento do ato de conversão de sociedade simples em sociedade empresária			
Data de encerramento do exercício social	31/12/2023		
Quantidade total de linhas do arquivo digital	46944		
	TERMO DE ENCERRAMENTO		
Nome Empresarial	PERFIL COMPUTACIONAL LTDA		
Natureza do Livro	LIVRO DIARIO		
Número de ordem	26		
Quantidade total de linhas do arquivo digital	46944		
Data de inicio	01/01/2023		
Data de término	31/12/2023		

Este documento é parte integrante de escrituração cuja autenticação se comprova pelo recibo de número E7.AC.9B.1B.84.ED.30.51.1D.92.93.ED.65.17.94.BA.99.95.65.1B-3, nos termos do Decreto nº 8.683/2016.

Este relatório foi gerado pelo Sistema Público de Escrituração Digital - Sped

Versão 10.2.1 do Visualizador

Página 1 de 1



MINISTÉRIO DA FAZENDA SECRETARIA DA RECEITA FEDERAL DO BRASIL SISTEMA PÚBLICO DE ESCRITURAÇÃO DIGITAL – Sped

Versão: 10.2.0

RECIBO DE ENTREGA DE ESCRITURAÇÃO CONTÁBIL DIGITAL

IDENTIFICAÇÃO DO TITULAR DA ESCRITURAÇÃO				
NIRE	CNPJ			
43203789542	02.543.216/0001-29			
NOME EMPRESARIAL				
PERFIL COMPUTACIONAL LTDA				

IDENTIFICAÇÃO DA ESCRITURAÇÃO	
FORMA DA ESCRITURAÇÃO CONTÁBIL	PERÍODO DA ESCRITURAÇÃO
Livro Diário (Completo - sem escrituração Auxiliar)	01/01/2023 a 31/12/2023
NATUREZA DO LIVRO	NÚMERO DO LIVRO
LIVRO DIARIO	26
IDENTIFICAÇÃO DO ARQUIVO (HASH)	
E7.AC.9B.1B.84.ED.30.51.1D.92.93.ED.65.17.94.BA.99.95.65.1B	

ESTE LIVRO FOI ASSINADO COM OS SEGUINTES CERTIFICADOS DIGITAIS:

QUALIFICAÇÃO DO SIGNATARIO	CPF/CNPJ	NOME	Nº SÉRIE DO CERTIFICADO	VALIDADE	RESPONSÁVEL LEGAL
Pessoa Jurídica (e-CNPJ ou e-PJ)	02543216000129	PERFIL COMPUTACIONAL LTDA:02543216000129	604594862258010616 7	09/02/2024 a 08/02/2025	Sim
Contador	49891570010	PAULO CESAR MOTTA:49891570010	619531105860947732 8	11/07/2023 a 10/07/2024	Não
Contador/Contabilista Responsável pelo Termo de Verificação para Fins de Substituição da ECD	49891570010	PAULO CESAR MOTTA:49891570010	619531105860947732 8	11/07/2023 a 10/07/2024	-

NÚMERO DO RECIBO:

E7.AC.9B.1B.84.ED.30.51.1D.92.93.ED .65.17.94.BA.99.95.65.1B-3

Escrituração recebida via Internet pelo Agente Receptor SERPRO					
em 07/05/2024 às 16:52:03					
D8.BC.A5.C9.3F.76.B1.37 BC.05.E6.79.FC.D5.75.FC					

Considera-se autenticado o livro contábil a que se refere este recibo, dispensando-se a autenticação de que trata o art. 39 da Lei nº 8.934/1994. Este recibo comprova a autenticação.

BASE LEGAL: Decreto nº 1.800/1996, com a alteração do Decreto nº 8.683/2016, e arts. 39, 39-A, 39-B da Lei nº 8.934/1994 com a alteração da Lei Complementar nº 1247/2014.

BALANÇO PATRIMONIAL

Entidade:

PERFIL COMPUTACIONAL LTDA

Período da Escrituração: 01/01/2023 a 31/12/2023

CNPJ: 02.543.216/0001-29

ls. 249

Número de Ordem do Livro: 26

Período Selecionado: 01 de Janeiro de 2023 a 31 de Dezembro de 2023

Descrição	Nota	Saldo Inicial	Saldo Final
ATIVO		R\$ 53.421.999,74	R\$ 76.557.226,67
ATIVO CIRCULANTE		R\$ 52.240.083,26	R\$ 76.124.210,53
DISPONIBILIDADES		R\$ 38.318.822,41	R\$ 55.168.916,76
BANCOS C/MOVIMENTO		R\$ 38.318.822,41	R\$ 55.168.916,76
BANCOS C/DISPOSICAO		R\$ 8.572.429,58	R\$ 5.536.023,93
BANCO DO BRASIL S/A - 2020-6		R\$ 6.883.314,84	R\$ 5.497.185,17
BANCO DO BRASIL S/A - 30185		R\$ 37.975,13	R\$ 0,00
BANCO DO BRASIL S/A - 30751-3		R\$ 0,00	R\$ 175,12
BANCO DO BRASIL S/A - 42163-4		R\$ 1.392.578,70	R\$ 6.317,02
BANCO DO BRASIL S/A - 42962-7		R\$ 31.169,36	R\$ 3.624,15
BANCO DO BRASIL S/A - 43945-2		R\$ 3.920,13	R\$ 2.048,42
BANCO DO BRASIL S/A - 5186-1		R\$ 223.471,42	R\$ 26.674,05
BANCOS C/APLICACOES FINANCEIRAS	3	R\$ 29.743.392,83	R\$ 49.627.892,83
BANCO DO BRASIL RF IND PRECO		R\$ 4.532.820,96	R\$ 4.532.820,96
BANCO DO BRASIL S/A C/APL.		R\$ 25.210.571,87	R\$ 45.095.071,87
TITULOS CAPITALIZACAO		R\$ 3.000,00	R\$ 5.000,00
BANCO DO BRASIL - CAPITALIZACAO		R\$ 3.000,00	R\$ 5.000,00
DIREITOS REALIZAVEIS		R\$ 13.921.260,85	R\$ 20.955.293,77
DUPLICATAS A RECEBER		R\$ 13.583.192,04	R\$ 20.289.124,42
CLIENTES NO PAIS		R\$ 13.583.192,04	R\$ 20.289.124,42
CLIENTES DIVERSOS		R\$ 13.583.192,04	R\$ 20.289.124,42
ADIANTAMENTO A TERCEIROS		R\$ 8.637,51	R\$ 40.209,75
ADIANTAMENTOS A TERCEIROS		R\$ 8.637,51	R\$ 40.209,75
ADIANTAMENTO DE FERIAS		R\$ 8.624,20	R\$ 40.194,84
ARREDONDAMENTOS DE SALARIO		R\$ 13,31	R\$ 14,91
TRIBUTOS A RECUPERAR		R\$ 270.879,27	R\$ 546.738,26
IMPOSTOS A RECUPERAR		R\$ 270.879,27	R\$ 546.738,26
COFINS A RECUPERAR		R\$ 0,00	R\$ 47.799,91
CSLL A RECUPERAR - CRF 4,65%		R\$ 148.997,35	R\$ 148.997,35
ICMS A RECUPERAR		R\$ 102.562,18	R\$ 319.831,10
ICMS ANTECIPADO A RECUPERAR		R\$ 0,00	R\$ 660,39
IRF A RECUPERAR		R\$ 19.319,74	R\$ 19.319,74
PIS A RECUPERAR		R\$ 0,00	R\$ 10.129,77
ESTOQUE		R\$ 10.052,03	R\$ 30.721,34

Este relatório foi gerado pelo Sistema Público de Escrituração Digital – Sped

Versão 10.2.0 do Visualizador

Página 1 de 3

Inserido ao protocolo **22.951.206-4** por: **Pedro Henrique Golin Linhares** em: 31/03/2025 12:49. A autenticidade deste documento pode ser validada no endereço: https://www.eprotocolo.pr.gov.br/spiweb/validarDocumento com o código: 1d3e5a42275b5642ba40f346676ed94.

BALANÇO PATRIMONIAL

Entidade:

PERFIL COMPUTACIONAL LTDA

Período da Escrituração: 01/01/2023 a 31/12/2023

CNPJ: 02.543.216/0001-29

is. <u>250</u>

Número de Ordem do Livro: 26

Período Selecionado: 01 de Janeiro de 2023 a 31 de Dezembro de 2023

Descrição	Nota	Saldo Inicial	Saldo Final
ESTOQUE		R\$ 10.052,03	R\$ 30.721,34
MERCADORIAS P/REVENDA		R\$ 10.052,03	R\$ 30.721,34
CREDITOS		R\$ 48.500,00	R\$ 48.500,00
OUTROS CREDITOS		R\$ 48.500,00	R\$ 48.500,00
LACAVE INDUSTRIA E COMERCIO DE BEBIDAS LTDA		R\$ 48.500,00	R\$ 48.500,00
ATIVO NÃO CIRCULANTE		R\$ 1.181.916,48	R\$ 433.016,14
ATIVO IMOBILIZADO		R\$ 1.181.916,48	R\$ 433.016,14
IMOBILIZADO		R\$ 1.181.916,48	R\$ 433.016,14
IMOBILIZACOES		R\$ 1.563.430,39	R\$ 879.554,49
EQUPTOS DE INFORMATICA		R\$ 689.803,75	R\$ 728.518,18
INSTALAÇÕES		R\$ 0,00	R\$ 168,67
INSTALAÇÕES TELEFONICA		R\$ 8.420,74	R\$ 8.420,74
MAQUINAS E EQUPTOS		R\$ 9.687,93	R\$ 9.687,93
MOVEIS E UTENSILIOS		R\$ 127.998,97	R\$ 132.758,97
VEICULOS		R\$ 727.519,00	R\$ 0,00
(-) (-)DEPREC.ACUMULADAS S/IMOBILIZADO		R\$ (381.513,91)	R\$ (446.538,35)
(-) DEPR.ACUM. S/VEICULOS		R\$ (35.202,85)	R\$ 0,00
(-) DEPR.ACUM.S/EQUPTOS DE INFORMATICA		R\$ (285.674,06)	R\$ (382.315,74)
(-) DEPR.ACUM.S/INST. TELEFONICA		R\$ (5.945,86)	R\$ (6.214,54)
(-) DEPR.ACUM.S/MOVEIS E UTENSILIOS		R\$ (54.691,14)	R\$ (58.008,07)
PASSIVO		R\$ 53.421.999,74	R\$ 76.557.226,67
PASSIVO CIRCULANTE		R\$ 26.534.846,64	R\$ 44.561.661,14
EXIGIVEL A CURTO PRAZO		R\$ 26.534.846,64	R\$ 44.561.661,14
DEBITOS P/DUPLS.A PAGAR		R\$ 23.739.410,33	R\$ 39.142.253,42
FORNECEDORES/CREDORES		R\$ 23.739.410,33	R\$ 39.142.253,42
FORNECEDORES DIVERSOS		R\$ 23.739.410,33	R\$ 39.142.253,42
DEBITOS TRIBUTARIOS E TRABALHISTA		R\$ 2.261.446,45	R\$ 4.575.713,31
DEBITOS TRABALHISTAS		R\$ 390.758,47	R\$ 969.757,44
FGTS A RECOLHER		R\$ 36.026,11	R\$ 82.825,13
INSS A RECOLHER		R\$ 102.388,36	R\$ 249.800,26
SALARIOS A PAGAR		R\$ 252.344,00	R\$ 637.132,05
DEBITOS TRIBUTARIOS		R\$ 1.870.687,98	R\$ 3.605.955,87
COFINS A RECOLHER		R\$ 64.501,67	R\$ 0,00

Este relatório foi gerado pelo Sistema Público de Escrituração Digital – Sped

Versão 10.2.0 do Visualizador

Página 2 de 3

BALANÇO PATRIMONIAL



Entidade:

PERFIL COMPUTACIONAL LTDA

Período da Escrituração: 01/01/2023 a 31/12/2023

CNPJ: 02.543.216/0001-29

Número de Ordem do Livro: 26

Período Selecionado: 01 de Janeiro de 2023 a 31 de Dezembro de 2023

Descrição	Nota	Saldo Inicial	Saldo Final
CRF RETIDO S/NFS A RECOLHER		R\$ 5.987,93	R\$ 44.571,47
CSLL A RECOLHER		R\$ 438.390,31	R\$ 626.324,92
ICMS A RECOLHER		R\$ 68.969,34	R\$ 239.434,87
ICMS A RECOLHER DIF ALIQTA N/C F/RS		R\$ 43.098,43	R\$ 201.734,62
ICMS ST A RECOLHER		R\$ 104,85	R\$ 104,85
IRF A RECOLHER		R\$ 155.095,75	R\$ 344.737,10
IRF RETIDO S/ NF		R\$ 61.315,67	R\$ 196.833,31
IRPJ A RECOLHER		R\$ 998.410,01	R\$ 1.938.518,37
ISSQN A RECOLHER		R\$ 21.808,11	R\$ 13.690,82
ISSQN RETIDO S/NFS		R\$ 5,26	R\$ 5,54
PIS A RECOLHER		R\$ 13.000,65	R\$ 0,00
DEBITOS		R\$ 533.989,86	R\$ 843.694,41
SOCIOS C/PARTICULAR		R\$ 1.113,71	R\$ 1.174,00
IGOR SIDNEI REOLON		R\$ 1.113,71	R\$ 1.174,00
PROVISÕES		R\$ 532.876,15	R\$ 842.520,41
ESTIMATIVA DE INSS S/RECIBO DE FÉRIAS		R\$ 717,84	R\$ 1.712,88
PROVISAO FGTS S/FERIAS		R\$ 31.655,84	R\$ 50.080,92
PROVISAO INSS S/FERIAS		R\$ 104.221,92	R\$ 164.656,55
PROVISAO P/PGTO FERIAS		R\$ 396.280,55	R\$ 626.070,06
PATRIMONIO LIQUIDO		R\$ 26.887.153,10	R\$ 31.995.565,53
CAPITAL		R\$ 15.000.000,00	R\$ 15.000.000,00
CAPITAL SOCIAL INTEGRALIZADO		R\$ 15.000.000,00	R\$ 15.000.000,00
CAPITAL SOCIAL INTEGRALIZADO		R\$ 15.000.000,00	R\$ 15.000.000,00
CAPITAL SOCIAL		R\$ 15.000.000,00	R\$ 15.000.000,00
RESERVAS		R\$ 11.887.153,10	R\$ 16.995.565,53
RESERVA DE CAPITAL		R\$ 11.887.153,10	R\$ 16.995.565,53
RESERVA DE CAPITAL		R\$ 11.887.153,10	R\$ 16.995.565,53
RESERVA DE LUCROS		R\$ 11.887.153,10	R\$ 16.995.565,53

Este relatório foi gerado pelo Sistema Público de Escrituração Digital - Sped

Versão 10.2.0 do Visualizador

Página 3 de 3

Inserido ao protocolo **22.951.206-4** por: **Pedro Henrique Golin Linhares** em: 31/03/2025 12:49. A autenticidade deste documento pode ser validada no endereço: https://www.eprotocolo.pr.gov.br/spiweb/validarDocumento com o código: 1d3e5a42275b5642ba40f346676ed94.

Entidade:

PERFIL COMPUTACIONAL LTDA

Período da Escrituração: 01/01/2023 a 31/12/2023

CNPJ: 02.543.216/0001-29

Número de Ordem do Livro: 26

Período Selecionado: 01 de Janeiro de 2023 a 31 de Dezembro de 2023

Descrição	Nota	Saldo anterior	Saldo atual
Receita Operacional		R\$ 112.165.251,05	R\$ 122.491.163,68
RECEITAS DE ALUGUÉIS		R\$ 0,00	R\$ 26.340,00
RECEITAS DE SERVIÇOS		R\$ 3.250.036,36	R\$ 2.747.938,97
REVENDAS DE MERCADORIAS		R\$ 108.915.123,69	R\$ 119.720.603,87
VENDAS DE SUCATAS		R\$ 91,00	R\$ 0,00
(-) (-) ICMS S/SUBST.TRIBUTARIA		R\$ (0,00)	R\$ (3.719,16)
(-) Deducoes Receita		R\$ (27.973.900,94)	R\$ (26.343.868,85)
(-) COFINS S/FATURAMENTO		R\$ (8.047.584,26)	R\$ (8.298.145,16)
(-) DEVOLUCOES DE VENDAS		R\$ (1.725.903,23)	R\$ (1.210.714,57)
(-) ICMS S/VENDAS		R\$ (14.945.301,44)	R\$ (11.135.833,15)
(-) ICMS S/VENDAS DIFERENCIAL CONSUMIDOR FINAL F/RS		R\$ (1.734.667,98)	R\$ (4.032.484,22)
(-) ISSQN S/SERVIÇOS		R\$ (75.565,46)	R\$ (65.437,83)
(-) PIS S/FATURAMENTO		R\$ (1.747.171,99)	R\$ (1.801.572,01)
(+) COFINS S/DEVOLUÇÃO DE VENDAS		R\$ 119.828,53	R\$ 87.142,69
(+) ICMS S/DEVOLUÇÃO DE VENDAS		R\$ 156.449,51	R\$ 94.256,26
(+) PIS S/DEVOLUÇÃO DE VENDAS		R\$ 26.015,38	R\$ 18.919,14
Receita Líquida		R\$ 84.191.350,11	R\$ 96.147.294,83
(-) Custos Mercadorias Vendidas		R\$ (58.644.037,93)	R\$ (65.617.087,10)
(-) COMPRA DE LICENCA DE SOFTWARE		R\$ (267.743,25)	R\$ (6.861,90)
(-) COMPRA DE MERCADORIAS		R\$ (66.505.080,81)	R\$ (62.566.087,84)
(-) CUSTO COM GARANTIA		R\$ (13.369.234,87)	R\$ (21.804.602,28)
(-) CUSTOS COM IMPORTAÇÃO		R\$ (1.315,91)	R\$ (0,00)
(-) FRETES S/COMPRAS		R\$ (7.408,02)	R\$ (61,10)
VARIAÇÃO DE ESTOQUE		R\$ (22.868,87)	R\$ 20.669,31
(-) COFINS S/COMPRAS		R\$ 6.091.269,05	R\$ 6.256.455,01
(-) DEVOLUÇÕES DE COMPRAS		R\$ 14.285,28	R\$ 5.461,45
(-) ICMS S/COMPRAS		R\$ 14.101.612,87	R\$ 11.119.630,94
(-) PIS S/COMPRAS		R\$ 1.322.446,60	R\$ 1.358.309,31
Lucro Bruto		R\$ 25.547.312,18	R\$ 30.530.207,73
(-) Despesas Com Pessoal - Vendas		R\$ (6.143.359,79)	R\$ (7.529.368,88)
(-) 13§ SALARIO		R\$ (262.789,73)	R\$ (459.341,16)
(-) COMISSÕES		R\$ (2.500.848,46)	R\$ (2.805.041,90)
(-) CONTRIBUICAO P/FGTS		R\$ (354.695,87)	R\$ (443.479,68)
(-) FERIAS		R\$ (418.979,42)	R\$ (665.506,29)
(-) HORAS EXTRAS		R\$ (761,01)	R\$ (354,33)
(-) INDENIZACOES TRABALHISTAS		R\$ (21.664,04)	R\$ (27.954,34)

Este relatório foi gerado pelo Sistema Público de Escrituração Digital - Sped

Versão 10.2.0 do Visualizador

Página 1 de 4



Entidade:

PERFIL COMPUTACIONAL LTDA

CNPJ: 02.543.216/0001-29

Número de Ordem do Livro: 26

Período da Escrituração: 01/01/2023 a 31/12/2023

Período Selecionado: 01 de Janeiro de 2023 a 31 de Dezembro de 2023

Descrição	Nota	Saldo anterior	Saldo atual
(-) PREVIDENCIA SOCIAL		R\$ (1.145.016,91)	R\$ (1.415.929,98)
(-) PRO-LABORE		R\$ (14.544,00)	R\$ (15.733,29)
(-) SALARIOS		R\$ (1.327.366,87)	R\$ (1.576.689,94)
(-) ALIMENTACAO		R\$ 137,00	R\$ (2,00)
(-) ASSISTENCIA MEDICA E SOCIAL		R\$ (96.830,48)	R\$ (118.900,97)
(-) VALE TRANSPORTE		R\$ (0,00)	R\$ (435,00)
(-) Despesas Operacionais		R\$ (1.510.416,87)	R\$ (1.677.290,49)
(-) DESPESAS C/AGUA		R\$ (7.629,63)	R\$ (8.334,24)
(-) DESPESAS C/ALIMENTACAO		R\$ (135.109,51)	R\$ (158.454,70)
(-) DESPESAS C/ALUGUEIS BEM MOVEIS		R\$ (1.552,62)	R\$ (0,00)
(-) DESPESAS C/ALUGUEIS DE IMOVEIS		R\$ (115.297,31)	R\$ (124.192,75)
(-) DESPESAS C/ANUIDADE		R\$ (4,00)	R\$ (0,00)
(-) DESPESAS C/ARMAZENAGEM		R\$ (8.490,12)	R\$ (2.265,18)
(-) DESPESAS C/ASSESSORIA		R\$ (56.416,95)	R\$ (72.999,19)
(-) DESPESAS C/BENS PERMANENTE		R\$ (1.572,00)	R\$ (2.337,91)
(-) DESPESAS C/BRINDES		R\$ (7.920,00)	R\$ (3.360,00)
(-) DESPESAS C/CARTORIO		R\$ (1.466,20)	R\$ (976,20)
(-) DESPESAS C/CERTIFICAÇÃO		R\$ (0,00)	R\$ (18.333,32)
(-) DESPESAS C/COLETA E ENTREGA		R\$ (4.621,88)	R\$ (1.235,00)
(-) DESPESAS C/COMBUSTIVEIS E LUBRIFICANTES		R\$ (14.796,33)	R\$ (15.997,72)
(-) DESPESAS C/COMEMORAÇÕES E FESTAS		R\$ (82.220,05)	R\$ (7.695,80)
(-) DESPESAS C/COMISSÕES		R\$ (5.000,02)	R\$ (7.488,00)
(-) DESPESAS C/COMUNICAÇÃO- TELEFONE		R\$ (14.672,19)	R\$ (16.446,19)
(-) DESPESAS C/CORREIOS E TELEGRAFOS		R\$ (2.382,72)	R\$ (5.730,11)
(-) DESPESAS C/CURSOS		R\$ (3.575,00)	R\$ (9.166,66)
(-) DESPESAS C/DEPRECIAÇÃO		R\$ (110.899,12)	R\$ (100.227,29)
(-) DESPESAS C/ENERGIA ELETRICA		R\$ (20.985,39)	R\$ (21.071,34)
(-) DESPESAS C/EVENTOS		R\$ (239,85)	R\$ (28.050,00)
(-) DESPESAS C/FEIRAS E EXPOSIÇÕES		R\$ (0,00)	R\$ (5.575,25)
(-) DESPESAS C/FRETES E DESPACHOS		R\$ (3.995,81)	R\$ (4.288,88)
(-) DESPESAS C/FRETES S/VENDAS		R\$ (19.709,26)	R\$ (6.484,87)
(-) DESPESAS C/GUINCHOS E GUINDASTES		R\$ (1.650,00)	R\$ (0,00)
(-) DESPESAS C/HONORARIOS ADVOCATICIOS		R\$ (115.422,31)	R\$ (3.500,00)
(-) DESPESAS C/HONORARIOS CONTABEIS		R\$ (120.914,56)	R\$ (126.150,32)

Este relatório foi gerado pelo Sistema Público de Escrituração Digital – Sped

Versão 10.2.0 do Visualizador

Página 2 de 4



Entidade:

PERFIL COMPUTACIONAL LTDA

Período da Escrituração: 01/01/2023 a 31/12/2023

CNPJ: 02.543.216/0001-29

Número de Ordem do Livro: 26

Período Selecionado: 01 de Janeiro de 2023 a 31 de Dezembro de 2023

Descrição	Nota	Saldo anterior	Saldo atual
(-) DESPESAS C/IMPRESSÕES		R\$ (20,00)	R\$ (20,00)
(-) DESPESAS C/LAVANDERIA		R\$ (0,00)	R\$ (180,00)
(-) DESPESAS C/MANUTENÇÃO E CONSERTOS		R\$ (20.630,35)	R\$ (78.902,82)
(-) DESPESAS C/MAT. EXPEDIENTE E ESCRITORIO		R\$ (19.426,55)	R\$ (8.813,45)
(-) DESPESAS C/MATERIAL DE CONSUMO		R\$ (2.431,61)	R\$ (3.023,95)
(-) DESPESAS C/MATERIAL DE LIMPEZA E COZINHA		R\$ (4.869,13)	R\$ (2.583,15)
(-) DESPESAS C/MENSALIDADES E CONTRIBUIÇÕES		R\$ (1.033,26)	R\$ (0,00)
(-) DESPESAS C/MONITORAMENTO		R\$ (3.153,72)	R\$ (3.325,56)
(-) DESPESAS C/PROC. DE DADOS		R\$ (53.670,97)	R\$ (76.730,17)
(-) DESPESAS C/PROPAGANDA E PUBLICIDADE		R\$ (285.445,00)	R\$ (231.147,06)
(-) DESPESAS C/REFORMAS		R\$ (0,00)	R\$ (100.645,17)
(-) DESPESAS C/SEGUROS		R\$ (16.665,45)	R\$ (11.982,88)
(-) DESPESAS C/VEICULOS		R\$ (4.711,00)	R\$ (35.780,00)
(-) DESPESAS C/VIAGENS E ESTADIAS		R\$ (241.817,00)	R\$ (373.795,36
(-) Despesas Tributarias		R\$ (60.561,39)	R\$ (1.153.783,84
(-) ICMS COMPETE - ES		R\$ (24.330,99)	R\$ (644.695,22)
(-) ICMS DIFERENCIAL DE ALÍQUOTA		R\$ (78,00)	R\$ (495.737,00)
(-) ICMS SUBST. TRIBUTARIA		R\$ (1.780,21)	R\$ (0,00)
(-) IMPOSTOS E TAXAS DIVERSAS		R\$ (29.845,76)	R\$ (10.497,87
(-) IOF		R\$ (3.494,40)	R\$ (1.777,76
(-) IPTU		R\$ (1.032,03)	R\$ (1.075,99)
(-) Despesas Financeiras		R\$ (1.490.670,84)	R\$ (1.377.832,44)
(-) JUROS S/CAPITAL PRÓPRIO		R\$ (1.469.025,65)	R\$ (1.276.002,29)
(-) JUROS PASSIVOS		R\$ (153,58)	R\$ (2.040,57)
(-) ACRESCIMOS LEGAIS S/IMPOSTOS		R\$ (12.131,62)	R\$ (80,40)
(-) DESCONTOS CONCEDIDOS		R\$ (3,86)	R\$ (434,47)
(-) DESPESAS BANCARIAS		R\$ (5.182,79)	R\$ (5.002,12)
(-) DESPESAS C/CARTAO DE CREDITO		R\$ (4.173,34)	R\$ (91.287,92
(-) DESPESAS C/SERVIÇOS DE COBRANÇAS		R\$ (0,00)	R\$ (2.984,67)
Receitas Financeiras		R\$ 2.347.847,85	R\$ 2.285.025,21
RENDIMENTOS S/APLIC.FINANCEIRAS		R\$ 2.457.854,79	R\$ 2.390.669,20
DESCONTOS OBTIDOS		R\$ 87,51	R\$ 400,39
JUROS ATIVOS		R\$ 4.195,79	R\$ 5.121,72
(-) COFINS S/RECEITA FINANCEIRA		R\$ (98.314,19)	R\$ (95.626,78)

Este relatório foi gerado pelo Sistema Público de Escrituração Digital - Sped

Versão 10.2.0 do Visualizador

Página 3 de 4



Entidade:

PERFIL COMPUTACIONAL LTDA

Período da Escrituração: 01/01/2023 a 31/12/2023

CNPJ: 02.543.216/0001-29

Número de Ordem do Livro: 26

Período Selecionado: 01 de Janeiro de 2023 a 31 de Dezembro de 2023

Descrição	Nota	Saldo anterior	Saldo atual
(-) PIS S/RECEITA FINANCEIRA		R\$ (15.976,05)	R\$ (15.539,32)
Resultado operacional líquido		R\$ 18.690.151,14	R\$ 21.076.957,29
(-) Despesas Não Operacionais		R\$ (283.235,38)	R\$ (35.635,15)
(-) MULTA TJ		R\$ (227.080,20)	R\$ (0,00)
(-) PERDAS BENS ATIVO IMOBILIZADO		R\$ (42.884,44)	R\$ (35.635,15)
(-) OUTRAS DESPESAS NÃO DEDUTÍVEIS		R\$ (13.270,74)	R\$ (0,00)
Receitas Não Operacionais		R\$ 1.877.991,05	R\$ 1.319.652,27
OUTRAS RECEITAS NAO OPERACIONAIS		R\$ 0,00	R\$ 168,71
RECUPERACAO DE DESPESA		R\$ 167.161,44	R\$ 22.715,31
RESSARCIMENTOS RECEBIDOS		R\$ 1.697.314,99	R\$ 1.215.367,25
VARIACAO MONETARIA ATIVA		R\$ 11.214,62	R\$ 0,00
GANHOS DE CAPITAL(LUCRO VENDA IMOBIL.)		R\$ 2.300,00	R\$ 81.401,00
Resultado Antes do IR		R\$ 20.284.906,81	R\$ 22.360.974,41
(-) Provisões		R\$ (6.857.404,07)	R\$ (7.570.966,52)
(-) PROVISÃO P/CONTR. SOCIAL		R\$ (1.824.078,19)	R\$ (2.013.482,72)
IRPJ DE SUBVENÇÕES GOVERNAMENTAIS PAT		R\$ 9.558,00	R\$ 11.523,75
(-) PROVISÃO P/IMPOSTO DE RENDA		R\$ (5.042.883,88)	R\$ (5.569.007,55)
LUCRO LÍQUIDO DO EXERCÍCIO		R\$ 13.427.502,74	R\$ 14.790.007,89

Este relatório foi gerado pelo Sistema Público de Escrituração Digital - Sped

Versão 10.2.0 do Visualizador

Página 4 de 4

		DEMONSTRAÇÃO	DE LUCROS	OU PREJUÍZOS ACU	MULADOS	Sped CONTABIL
Entidade:	PERFIL CO	MPUTACIONAL LTDA				
Período da Escrituração:	01/01/2023	a 31/12/2023	CNP	02.543.216/0001-29	Número de Ordem do Livro: 26	
Período Selecionado:	01 de Janei	ro de 2023 a 31 de Dez	embro de 2023			
				Código de Aglutinação das Contas	de Patrimônio Líquido	
Histórico				RESULTADO DO EXER	RCICIO (R\$)	
Saldo Anterior de Lucros Acumulados						0,00
Ajustes Credores de Períodos-base Anterior	es					0,00
Reversão de Reservas						0,00
Outros Recursos						0,00
Lucro Líquido do Ano						24.403.580,14
(-)Saldo Anterior de Prejuízo Acumulados						0,00
(-)Ajustes Devedores de Períodos-base Ante	eriores					(-)9.681.595,46
(-)Prejuízo Líquido do Ano						(-)9.613.572,25
Transferências para Reservas	o					(-)5.108.412,43
Dividendos ou Lucros Distribuídos, Pagos ou	u Creditados					0,00
Parceia dos Lucros Incorporados ao Capital						0,00
						0,00
						0,00
						0,00
Notas						0,00

Este relatório foi gerado pelo Sistema Público de Escrituração Digital - Sped

10.2.0

Inserido ao protocolo 22.951.206-4 por: Pedro Henrique Golin Linhares em: 31/03/2025 12:49. A autenticidade deste documento pode ser validada no endereço: https://www.eprotocolo.pr.gov.br/spiweb/validarDocumento com o código: 1d3e5a42275b5642ba40f346676ed94.





PODER JUDICIÁRIO TRIBUNAL DE JUSTIÇA DO ESTADO DO ESPÍRITO SANTO R. Des. Homero Mafra, 60 Enseada do Suá, Vitória - ES | CEP: 29.050-275 | Tel: (27) 3334-2000.

CERTIDÃO NEGATIVA DE PRIMEIRA INSTÂNCIA NATUREZA <u>DE RECUPERAÇÃO JUDICIAL E EXTRAJUDICIAL (FALÊNCIA E</u> <u>CONCORDATA)</u>

Dados da Certidão

Razão Social: PERFIL COMPUTACIONAL LTDA

CNPJ: Data de Expedição: Nº da Certidão: ENDEREÇO	02.543.216/0011-09 18/03/2025 13:41:50 * 2024419320 *	Validade:	30 DIAS
Município: Logradouro: Complemento: CONTATO	SERRA AV. ACESSO RODOVIÁRIO QUADRA 06 LM 01 QUADRA 01 L-M18 AM23 SALA 107	Bairro: Número: CEP:	INTERMODAL DA SERRA SN 29.161-376
Email:	PERFIL@PERFIL.INF.BR	Telefone Fixo: Telefone Celular:	(54) 2628-8300 - <i>NÃO INFORMADO -</i>

CERTIFICA que, consultando a base de dados do Sistema de Gerenciamento de Processos do Poder Judiciário do Estado do Espírito Santo (E-Jud, SIEP, PROJUDI e PJe) até a presente data e hora, **NADA CONSTA** contra o solicitante .

- Observações

- a. Certidão expedida gratuitamente através da Internet;
- b. Os dados do(a) solicitante acima informados são de sua responsabilidade, devendo a titularidade ser conferida pelo interessado e/ou destinatário;
- c. O prazo de validade desta certidão é de 30 (trinta) dias, contados da data da expedição, conforme disposto no art. 467 do Código de Normas da Corregedoria Geral da Justiça. Após essa data será necessária a emissão de uma nova certidão;
- d. A autenticidade desta certidão poderá ser confirmada na página do Tribunal de Justiça do Estado do Espírito Santo www.tjes.jus.br -, utilizando o número da certidão acima identificado;
- e. Em relação as comarcas da entrância especial (Vitória/Vila Velha/Cariacica/Serra/Viana), as ações de: execução fiscal estadual, falência e recuperação judicial, e auditoria militar, tramitam, apenas, no juízo de Vitória;
- f. As ações de natureza cível abrangem inclusive aquelas que tramitam nas varas de Órfãos e Sucessões (Tutela, Curatela, Interdição,...), Juizado Especial Cível, Juizado Especial da Fazenda Pública, Execução Fiscal e Execução Patrimonial (observado o item **e**);
- g. As ações de natureza criminal abrangem, dentre outras: as de auditoria militar e de juizados especiais criminais;

h. As matérias atinentes as varas de família e infância e juventude são objeto de certidão específica;

 i. A base de dados do sistema de gerenciamento processual (1ª INSTÂNCIA: eJUD, SIEP, PROJUDI, PJe-1G; 2ª INSTÂNCIA: Sistema de Segunda Instância, PJe-2G) contém o registro de todos os processos distribuídos no Judiciário do Estado do Espírito Santo, com exceção do SEEU;

j. A certidão negativa referente ao Sistema Eletrônico de Execução Unificado – SEEU deverá ser requerida ao Cartório do Ofício de Distribuidor da Comarca, conforme Ato Normativo Conjunto nº. 009/2021.





CERTIDÃO JUDICIAL CÍVEL NEGATIVA

Farroupilha, 03 de fevereiro de 2025, às 13h08min

Assinado eletronicamente por Rio Grande Do Sul Poder Judiciario Confira autenticidade em https://www.tjrs.jus.br/verificadocs, informando 0001551910104.

Página 1/2



ESTADO DO RIO GRANDE DO SUL PODER JUDICIÁRIO TRIBUNAL DE JUSTIÇA



DOCUMENTO ASSINADO POR

RIO GRANDE DO SUL PODER JUDICIARIO

DATA 03/02/2025 13h08min



Este é um documento eletrônico assinado digitalmente conforme Lei Federal nº 11.419/2006 de 19/12/2006, art. 1º, parágrafo 2º, inciso III.

Para conferência do conteúdo deste documento, acesse, na internet, o endereço https://www.tjrs.jus.br/verificadocs e digite o seguinte

número verificador: 0001551918104



Página 2/2

Inserido ao protocolo **22.951.206-4** por: **Pedro Henrique Golin Linhares** em: 31/03/2025 12:49. A autenticidade deste documento pode ser validada no endereço: https://www.eprotocolo.pr.gov.br/spiweb/validarDocumento com o código: **1d3e5a42275b5642ba40f346676ed94**.





MINISTÉRIO DA FAZENDA Secretaria da Receita Federal do Brasil Procuradoria-Geral da Fazenda Nacional

CERTIDÃO NEGATIVA DE DÉBITOS RELATIVOS AOS TRIBUTOS FEDERAIS E À DÍVIDA ATIVA DA UNIÃO

Nome: PERFIL COMPUTACIONAL LTDA CNPJ: 02.543.216/0001-29

Ressalvado o direito de a Fazenda Nacional cobrar e inscrever quaisquer dívidas de responsabilidade do sujeito passivo acima identificado que vierem a ser apuradas, é certificado que não constam pendências em seu nome, relativas a créditos tributários administrados pela Secretaria da Receita Federal do Brasil (RFB) e a inscrições em Dívida Ativa da União (DAU) junto à Procuradoria-Geral da Fazenda Nacional (PGFN).

Esta certidão é válida para o estabelecimento matriz e suas filiais e, no caso de ente federativo, para todos os órgãos e fundos públicos da administração direta a ele vinculados. Refere-se à situação do sujeito passivo no âmbito da RFB e da PGFN e abrange inclusive as contribuições sociais previstas nas alíneas 'a' a 'd' do parágrafo único do art. 11 da Lei n^o 8.212, de 24 de julho de 1991.

A aceitação desta certidão está condicionada à verificação de sua autenticidade na Internet, nos endereços <http://rfb.gov.br> ou <http://www.pgfn.gov.br>.

Certidão emitida gratuitamente com base na Portaria Conjunta RFB/PGFN n^o 1.751, de 2/10/2014. Emitida às 10:22:54 do dia 03/02/2025 <hora e data de Brasília>. Válida até 02/08/2025. Código de controle da certidão: **6E04.1416.F7FE.A3E5** Qualquer rasura ou emenda invalidará este documento.





ESTADO DO ESPÍRITO SANTO SECRETARIA DE ESTADO DA FAZENDA

Certidão Negativa de Débitos para com a Fazenda Pública Estadual - MOD. 2

Certidão Nº 20250000348604

Identificação do Requerente: CNPJ Nº 02.543.216/0011-09

Certificamos que, até a presente data, não existe débito contra o portador do Cadastro de Pessoa Jurídica acima especificado, ficando ressalvada à Fazenda Pública Estadual o direito de cobrar quaisquer dívidas que venham a ser apuradas.

Certidão emitida via Sistema Eletrônico de Processamento de Dados, nos termos do Regulamento do ICMS/ES, aprovado pelo Decreto nº 1.090-R, de 25 de outubro de 2002.

Certidão emitida em 25/02/2025, válida até 26/05/2025.

A autenticidade deste documento poderá ser confirmada via internet por meio do endereço **www.sefaz.es.gov.br** ou em qualquer Agência da Receita Estadual.

Vitória, 25/02/2025.

Autenticação eletrônica: 000A.DA3D.1710.CBAF





PREFEITURA MUNICIPAL DA SERRA Secretaria Municipal da Fazenda



CEP: 29176-439 - RUA MAESTRO ANTONIO CICERO, Nº 111 CACAROCA Serra ES

Certidão Negativa de Débitos

N° da Certidão: 11971846/2025

Data Geração:10/03/2025

Data Validade: 10/05/2025

CERTIFICAMOS, que não constam em nome do sujeito passivo identificado, nesta data, débitos com a Fazenda Pública Municipal, ressalvando o direito do município de cobrar quaisquer débitos que vierem a ser conhecidos e apurados após a expedição.

Essa certidão com base no art. 178 da Lei 3833, de 29 de dezembro de 2011, Código Tributário Municipal (CTM).

		110		~
10	<u>or</u>	1111	$\sim \sim$	$\sim \sim \sim$
	—		L.A	
10	<u> </u>		20	çuo

Ccm 4703316 InscrMunicipal 4703316 Situação: Ativo Razão Social PERFIL COMPUTACIONAL LTDA CNPJ / CPF 02.543.216/0011-09 Inscrição Estadual/RG Endereço 29161-376 - AVENIDA ACESSO RODOVIARIO, QUADRA 06 LM 01;QUADRA 01 L-M18 a Bairro TERMINAL INTERMODAL DA SERRA Cidade SERRA Estado ES	
Serra, Segunda-feira, 10 de Março de 2025 N° da Certidão: 11971846/2025 Inscrição: 4703316	
Tanto a veracidade da informação quanto a manutenção da condição de não devedor poderá ser verificada na seguinte página da Internet: http://www.serra.es.gov.br/	
ATENÇÃO: Qualquer rasura ou emenda INVALIDARÁ este documento.	

Inserido ao protocolo **22.951.206-4** por: **Pedro Henrique Golin Linhares** em: 31/03/2025 12:49. A autenticidade deste documento pode ser validada no endereço: https://www.eprotocolo.pr.gov.br/spiweb/validarDocumento com o código: 1d3e5a42275b5642ba40f346676ed94.



Página 1 de 1



PODER JUDICIÁRIO JUSTIÇA DO TRABALHO

CERTIDÃO NEGATIVA DE DÉBITOS TRABALHISTAS

Nome: PERFIL COMPUTACIONAL LTDA (MATRIZ E FILIAIS) CNPJ: 02.543.216/0011-09 Certidão n°: 86456606/2024 Expedição: 16/12/2024, às 09:12:15 Validade: 14/06/2025 - 180 (cento e oitenta) dias, contados da data de sua expedição.

Certifica-se que **PERFIL COMPUTACIONAL LTDA (MATRIZ E FILIAIS)**, inscrito(a) no CNPJ sob o n° **02.543.216/0011-09**, **NÃO CONSTA** como inadimplente no Banco Nacional de Devedores Trabalhistas. Certidão emitida com base nos arts. 642-A e 883-A da Consolidação das Leis do Trabalho, acrescentados pelas Leis ns.º 12.440/2011 e

13.467/2017, e no Ato 01/2022 da CGJT, de 21 de janeiro de 2022. Os dados constantes desta Certidão são de responsabilidade dos Tribunais do Trabalho.

No caso de pessoa jurídica, a Certidão atesta a empresa em relação a todos os seus estabelecimentos, agências ou filiais.

A aceitação desta certidão condiciona-se à verificação de sua autenticidade no portal do Tribunal Superior do Trabalho na Internet (http://www.tst.jus.br). Certidão emitida gratuitamente.

INFORMAÇÃO IMPORTANTE

Do Banco Nacional de Devedores Trabalhistas constam os dados necessários à identificação das pessoas naturais e jurídicas inadimplentes perante a Justiça do Trabalho quanto às obrigações estabelecidas em sentença condenatória transitada em julgado ou em acordos judiciais trabalhistas, inclusive no concernente aos recolhimentos previdenciários, a honorários, a custas, a emolumentos ou a recolhimentos determinados em lei; ou decorrentes de execução de acordos firmados perante o Ministério Público do Trabalho, Comissão de Conciliação Prévia ou demais títulos que, por disposição legal, contiver força executiva.

CNH Digital

Departamento Nacional de Trânsito





QR-CODE



Documento assinado com certificado digital em conformidade com a Medida Provisória nº 2200-2/2001. Sua validade poderá ser confirmada por meio da comparação deste arquivo digital com o arquivo de assinatura (.p7s) no endereço: < http://www.serpro.gov.br/assinador-digital >.

SERPRO / DENATRAN

CNH Digital

Departamento Nacional de Trânsito





QR-CODE



Documento assinado com certificado digital em conformidade com a Medida Provisória nº 2200-2/2001. Sua validade poderá ser confirmada por meio do programa Assinador Serpro.

As orientações para instalar o Assinador Serpro e realizar a validação do documento digital estão disponíveis em: < http://www.serpro.gov.br/assinador-digital >, opção Validar Assinatura.

SERPRO / DENATRAN

C/	REPÚBLICA FEDEF	RATIVA DO) BRASIL DA JURÍDIO	CA	
NÚMERO DE INSCRIÇÃO 02.543.216/0011-09 FILIAL	COMPROVANTE DE INS CADA	CRIÇÃO E DE STRAL	SITUAÇÃO	DATA DE ABERTURA 03/03/2020	
NOME EMPRESARIAL PERFIL COMPUTACIONAL L	TDA				
TÍTULO DO ESTABELECIMENTO (NON *******	E DE FANTASIA)				PORTE DEMAIS
CÓDIGO E DESCRIÇÃO DA ATIVIDADE 46.51-6-01 - Comércio atacad	ECONÔMICA PRINCIPAL lista de equipamentos de inform	iática			
CÓDIGO E DESCRIÇÃO DAS ATIVIDAE 46.51-6-02 - Comércio atacad 46.52-4-00 - Comércio atacad 62.04-0-00 - Consultoria em t 62.09-1-00 - Suporte técnico, 74.90-1-04 - Atividades de int 77.33-1-00 - Aluguel de máqu	DES ECONÔMICAS SECUNDÁRIAS lista de suprimentos para inform lista de componentes eletrônico ecnologia da informação manutenção e outros serviços d rermediação e agenciamento de linas e equipamentos para escri	nática s e equipamento em tecnologia da serviços e negóo tórios	s de telefonia e informação cios em geral, e	comunicação xceto imobiliários	
CÓDIGO E DESCRIÇÃO DA NATUREZA 206-2 - Sociedade Empresári	a Jurídica a Limitada				
LOGRADOURO AV ACESSO RODOVIARIO		NÚMERO S/N	COMPLEMENTO QUADRA06 L M23 SALA 10	M 01 QUADRA01 L 7	-M18 A
CEP 29.161-376 SER	RO/DISTRITO MINAL INTERMODAL DA RA	MUNICÍPIO SERRA			UF ES
ENDEREÇO ELETRÔNICO IGOR@PERFIL.INF.BR		TELEFONE (54) 2628-832	1		
ENTE FEDERATIVO RESPONSÁVEL (E *****	FR)				
SITUAÇÃO CADASTRAL ATIVA			DA 03	TA DA SITUAÇÃO CADAS / 03/2020	TRAL
MOTIVO DE SITUAÇÃO CADASTRAL					
SITUAÇÃO ESPECIAL *******			DA ***	TA DA SITUAÇÃO ESPEC	IAL
L					

Aprovado pela Instrução Normativa RFB nº 2.119, de 06 de dezembro de 2022.

Emitido no dia 12/02/2025 às 10:27:21 (data e hora de Brasília).

Página: 1/1

FIS. <u>266</u>

Empresa:	PERFIL COMPUTACIONAL LTDA
Inscrição:	02.543.216/0001-29
Período:	01/01/2023 - 31/12/2023

0001

0008

OTOCO

Fls. <u>267</u>

Mov. <u>31</u>

COEFICIENTES DE ANÁLISES EM 31/12/2023

Coeficiente	Fórmula	Valor	Resultado
Índice de Liquidez Geral	Ativo Circulante + Realizável Longo Prazo	76.124.210,53 + 0,00	1,71
	Passivo Circulante + Passivo Não-Circulante	44.561.661,14 + 0,00	
Índice de Liquidez Corrente	Ativo Circulante	76.124.210,53	1,71
	Passivo Circulante	44.561.661,14	
Índice de Liquidez Seca	Ativo Circulante - Estoque	76.124.210,53 - 30.721,34	1,71
	Passivo Circulante	44.561.661,14	
Índice de Solvência Geral	Ativo	76.557.226,67	1,72
	Passivo Circulante + Passivo Não-Circulante	44.561.661,14 + 0,00	
Grau de Endividamento	Passivo Circulante + Passivo Não-Circulante	44.561.661,14 + 0,00	0,58
	Ativo	76.557.226,67	

IGOR SIDNEI REOLON:805 12771049

Dados: 2024.06.18 09:38:43 -03'00'

IGOR SIDNEI REOLON Admnistrador 805.127.710-49

PAULO CESAR MOTTA:4989157 0010 Assinado de forma digital por PAULO CESAR MOTTA:49891570010 Dados: 2024.06.18 09:22:24 0010 -03'00' PAULO CESAR MOTTA Reg. no CRC - RS sob o No. 61363 CPF: 498.915.700-10



SINTEGRA/ICMS Consulta Pública ao Cadastro Estado do Espírito Santo





Cadastro atualizado até: 12/02/2025

IDENTIFICAÇÃO - PESSOA JURÍDICA

 CNPJ:
 02543216001109
 Inscrição Estadual:
 083.644.61-0

 Razão
 PERFIL COMPUTACIONAL LTDA

 Social:
 Inscrição Estadual:
 083.644.61-0

ENDEREÇO

Logradouro:	AVENIDA ACESSO RODOVIA	ARIO	
Número:	SN	Complemento:	QUADRA 6, L-M01, QUADRA 01, L-M18 A M23, SALA 107
Bairro:	TERMINAL INTERMODAL DA	SERRA	
Município:	SERRA	UF:	ES
CEP:	29161376	Telefone:	(0054) 30281551

INFORMAÇÕES COMPLEMENTARES

Atividade Econômica:

COMERCIO ATACADISTA DE EQUIPAMENTOS DE INFORMATICA

COMERCIO ATACADISTA DE SUPRIMENTOS PARA INFORMATICA

COM ATACAD DE COMPONTENTES ELETRONICOS E EQUIP DE TELEFONIA E COMUNICA

CONSULTORIA EM TECNOLOGIA DA INFORMACAO

SUPORTE TECNICO, MANUTENCAO E OUTROS SERVICOS EM TECNOLOGIA DA INFORMAC

ATIV INTERMEDIACAO AGENC SERV/NEGOCIOS EM GERAL, EXC IMOBILIARIOS

ALUGUEL DEMAQUINAS E EQUIPAMENTOS PARA ESCRITORIO

Data de Inicio de Atividade:	11/03/2020
Situação Cadastral Vigente:	HABILITADO
Data desta Situação Cadastral:	11/03/2020
Regime de Apuração:	ORDINÁRIO
Devedor contumaz:	Não

A Inscrição Estadual (IE) com situação cadastral vigente **HABILITADO** indica que a empresa está **APTA** a realizar operações como contribuinte do ICMS.

Já a IE com situação cadastral **NÃO HABILITADA** indica que a empresa **NÃO** está **APTA** a realizar operações como contribuinte do ICMS, <u>caso mantenha entre as suas atividades pelo</u> <u>menos um CNAE cuja inscrição estadual seja obrigatória</u>. Caso a empresa não pertença a um CNAE cuja inscrição seja obrigatória e o CNPJ esteja ATIVO (consultar o site da Receita Federal do Brasil ? http://www.receita.fazenda.gov.br), a empresa poderá ser destinatária de mercadorias, bens e serviços <u>como CONSUMIDOR FINAL</u>. Neste caso, <u>o número da Inscrição Estadual NÃO deverá constar em documentos que acobertem operações tributáveis pelo ICMS</u>.

A lista dos CNAEs obrigados à Inscrição Estadual está disponível no endereço: ftp://ftp.sefaz.es.gov.br/CNAE-F/cnaes_obrigadas_a_inscricao.pdf

OBSERVAÇÃO: Os dados acima são baseados em informações fornecidas pelo contribuinte, estando sujeitos a posterior confirmação pelo Fisco.

VOLTAR



© Copyright 2003/2025 **Secretaria de Estado da Fazenda do Espírito Santo** Av. João Batista Parra . nº600 . Ed. Aureliano Hoffman . Enseada do Suá . Vitória-ES CEP: 29050-375 . CNPJ: 27.080.571/0001-30



PREFEITURA MUNICIPAL DA SERRA

RUA MAESTRO ANTONIO CICERO, 111 CEP 29176-439 - CACAROCA - Serra - ES



Cadastro - OCM Inscrição Municipal CPF(CNP) Inscrição Estadual Data Inicio Aludade AT03316 AT03316 Q2.543.216/0011-09 Inscrição Estadual Data Inicio Aludade Nome PERFIL COMPUTACIONAL LTDA CEP Vome Fantasia Cidades 29161-376 AVENIDA ACESSO RODOVIARIO, QUADRA 06 LM 01;QUADRA 01 L-M18 a 29161-377 Bairo Cidades SERRA UF ES Silvação Cudadital Simples Nacional Tipo ISS Ativo Qata de Enceramento Com Nao Nao Tipo ISS Ativo Oata de Enceramento Com Código e descrição da atividade condinica principal 4 - ISS Variável Codigo e descrição da atividade de licença 4652-4/00 - Comércio atacadista de componentes eletrônicos e equipamentos de telefonia e comunicação Código e descrição das atividades de licença 4652-4/00 - Comércio atacadista de componentes eletrônicos e equipamentos de telefonia e comunicação Código e descrição das atividades de licença 4652-4/00 - Comércio atacadista de componentes eletrônicos e equipamentos de telefonia e comunicação Data Enceramento Código e descrição das atividades de licença Gégo e descrição das atividades de licença Código e descrição das atividades de licenção		Comprov	vante de ins	criçao e d	e situação ca	dastral	
Vome PERFIL COMPUTACIONAL LTDA Vome Fantasia Enderaço 29161-37 - AVENIDA ACESSO RODOVIARIO, QUADRA 06 LM 01;QUADRA 01 L-M18 a 29161-37 EP TERMINAL INTERMODAL DA Cidade SERRA Cidade Straga Cidade Straga Nao Cidade Straga Nao Cidade Straga Nao Cidade Straga Cidade Straga Nao Cidade A Straga Cidade Straga Cidade Straga Cidade Straga Cidade Cinter Cidade Cinter Cinter Cidade Cinter Cidade Cinter	Cadastro - CCM 4703316	Inscrição Municipa 4703316	al CPF/CNPJ 6 02.543.2	216/0011-09	Inscrição Estadual	Data Ini	cio Atividade 03/03/2020
Ideme Fantasia Indeveço 19161-376 - AVENIDA ACESSO RODOVIARIO, QUADRA 06 LM 01;QUADRA 01 L-M18 a 29161-37 ETRINIAL INTERMODAL DA Cidade SERA Cidade SERA Cidade Catastral Simples Nacional Nao Cidade Catastral Nao Cidade Catastral Nao Cidade Catastral Cidade Cidade Comércio atacadista de equipamentos de informática Cidago e descrição da atividades de licença Cidago e descrição das atividades de serviço Cidago e descrição das atividades de serviço Cidago e descrição das atividades de licença Cidago e descrição das atividades de licença Cidago e descrição das atividades de licença Cidago e descrição das atividades de serviço Cidago e descrição das atividades de serviço Cidago e descrição das atividades de serviço Cidago e descrição das atividades de licença Cidago e descrição das atividades de serviço Cidago e descrição das atividades de licença Cidago e descrição das atividades de serviço Cidago e descrição das atividades de licença Cidago e descrição das atividades de serviço Cidago e descrição das atividades de licença Cidago e descrição componentes eletrônicos e equipamentos de telefonia e comunicação Cidago e descrição componentes eletrô	Iome PERFIL COMPU	TACIONAL LTD	Α				
Indereço 19161-376 - AVENIDA ACESSO RODOVIARIO, QUADRA 06 LM 01;QUADRA 01 L-M18 a CEP 29161-371 Iairo CLidade SERRA U LF ES Ituação Cadastral Simples Nacional A SERRA Data de Encerramento Com Nao 4 - ISS Variável Data de Encerramento Com Vidade Livre Vidade Livre Vidade Livre Vidago e descrição da atividade econômica principal 1651-6/01 - Comércio atacadista de equipamentos de Informática 2020 e descrição da atividade de licença 1652-4/00 - Comércio atacadista de componentes eletrônicos e equipamentos de telefonia e comunicação 2020 e descrição das atividades de licença 1652-4/00 - Comércio atacadista de componentes eletrônicos e equipamentos de telefonia e comunicação 2020 o descrição das atividades de licença 1652-4/00 - Comércio atacadista de componentes eletrônicos e equipamentos de telefonia e comunicação 2020 o descrição das atividades de licença 1652-4/00 - Comércio atacadista de componentes eletrônicos e equipamentos de telefonia e comunicação 2020 o descrição das atividades de serviço 2020 - Comércio atacadista de componentes eletrônicos e equipamentos de telefonia e comunicação 2020 o Comércio atacadista de equipamentos de informática 1651-6/01 - Comércio atacadista de suprimentos para informática 1652-4/00 - Comércio atacadista de suprimentos para informática 1652-100 - Comércio atacadista de componentes eletrônicos e equipamentos de 1209-1/00 - Suporte técnico, manutenção e agenciamento de serviços e negócios em 1209-1/00 - Suporte técnico, manutenção e agenciamento de serviços e negócios em 1209-1/04 - Atividades de intermediação e agenciamento de serviços e negócios em 1209-1/04 - Atuguel de máquinas e equipamentos para escritórios 1209-1/04 - Atuguel de máquinas e equipamentos para escritórios	Iome Fantasia						
Sarro Citade UF FERMINAL INTERMODAL DA SERRA UF Simples Nacional Tipo 155 Data de Encerramento Com Ativo Nao 4 - 1SS Variável Data de Encerramento Com Vividade Livre 200 200 200 Data de Encerramento Com 200 comércio atacadista de equipamentos de Informática 200 200 200 comércio atacadista de equipamentos de Informática 200 200 200 comércio atacadista de componentes eletrônicos e equipamentos de telefonia e comunicação 200 200 comércio atacadista de componentes eletrônicos e equipamentos de telefonia e comunicação 200 200 comércio atacadista de componentes eletrônicos e equipamentos de telefonia e comunicação 200 200 comércio atacadista de suprimentos para informática 1651-6/01 - Comércio atacadista de suprimentos para informática 1651-6/01 - Comércio atacadista de componentes eletrônicos e equipamentos de elefonia e comunicação 200 200 200 comércio atacadista de suprimentos para informática 1652-4/00 - Consultoria em tecnología da informação 200 200 1/00 - Suporte técnico, manutenção e outros serviços e metecnología da informação 200 <t< td=""><td>Endereço 29161-376 - AVE</td><th>NIDA ACESSO</th><td>RODOVIARIO,</td><td>QUADRA 0</td><th>6 LM 01;QUADRA 0</th><td>1 L-M18 a</td><td>CEP 29161-376</td></t<>	Endereço 29161-376 - AVE	NIDA ACESSO	RODOVIARIO,	QUADRA 0	6 LM 01;QUADRA 0	1 L-M18 a	CEP 29161-376
Bitugção Cadastral Simples Nacional Tipo ISS Data de Encerramento Com Vidade Livre Código e descrição da atividade econômica principal K651-6/01 - Comércio atacadista de equipamentos de informática Código e descrição da atividades de licença K652-4/00 - Comércio atacadista de componentes eletrônicos e equipamentos de telefonia e comunicação Código e descrição das atividades de licença K652-4/00 - Comércio atacadista de componentes eletrônicos e equipamentos de telefonia e comunicação Código e descrição das atividades de licença K652-4/00 - Comércio atacadista de componentes eletrônicos e equipamentos de telefonia e comunicação Código e descrição das atividades de serviço Código e descrição das atividades de servição Código e descrição atacadista de componentes eletrônicos e equipamentos de elefonia e comunicação Código e descrição e agenciamento de serviços e negócios em geral, exceto imobilitários Codigo e deseserviço e a	Bairro FERMINAL INTE	RMODAL DA	Cldade SERRA				U.F ES
Dádigo e descrição da atividade econômica principal 4651-6/01 - Comércio atacadista de equipamentos de informática 2ódigo e descrição da atividade de licença 4652-4/00 - Comércio atacadista de componentes eletrônicos e equipamentos de telefonia e comunicação 2ódigo e descrição das atividades de licença 4652-4/00 - Comércio atacadista de componentes eletrônicos e equipamentos de telefonia e comunicação 2ódigo e descrição das atividades de licença 4652-4/00 - Comércio atacadista de componentes eletrônicos e equipamentos de telefonia e comunicação 2ódigo e descrição das atividades de serviço 2ódigo e descrição dos CNAEs 10616-6/01 - Comércio atacadista de equipamentos de informática 1651-6/01 - Comércio atacadista de suprimentos para informática 1652-4/00 - Consultoria em tecnologia da informação 3204-0/00 - Consultoria em tecnologia da informação 3209-1/00 - Suporte técnico, manutenção e outros serviços em tecnologia da informação 1733-1/00 - Aluguel de máquinas e equipamentos para escritórios veracidade da informação poderá ser verificada na seguinte página da Internet: https://www.serra.es.gov have de Verificação:	Situação Cadastral Ativo	Sir Na	nples Nacional 10	Tip 4 -	o ISS ISS Variável	Data de Encer	ramento Ccm
bidigo e descrição da atividade econômica principal 1651-6/01 - Comércio atacadista de equipamentos de informática 2001go e descrição da atividade de licença 1652-4/00 - Comércio atacadista de componentes eletrônicos e equipamentos de telefonia e comunicação 2001go e descrição das atividades de licença 1652-4/00 - Comércio atacadista de componentes eletrônicos e equipamentos de telefonia e comunicação 2001go e descrição das atividades de serviço 2011go e descrição das atividades de serviço 2012 e descrição das atividades de serviço 2013 e descrição das atividades de serviço 2014 e comúnicação 2020 e descrição confércio atacadista de componentes eletrônicos e equipamentos de elefonia e comunicação 2020 e dovo Comércio atacadista de componentes eletrônicos e equipamentos de elefonia e comunicação 2020 e dovo Consultoria em tecnologia da informação 2020 e dovo Suporte técnico, manutenção e agenciamento de serviços e negócios em ge							
Código e descrição da atividade de licença 4652-4/00 - Comércio atacadista de componentes eletrônicos e equipamentos de telefonia e comunicação Código e descrição das atividades de licença 4852-4/00 - Comércio atacadista de componentes eletrônicos e equipamentos de telefonia e comunicação Código e descrição das atividades de serviço Código e descrição das atividades de serviço Código e descrição das atividades de serviço Código e descrição das cNAEs 1651-6/01 - Comércio atacadista de equipamentos para informática 1652-4/00 - Comércio atacadista de componentes eletrônicos e equipamentos de elefonia e comunicação 1652-4/00 - Comércio atacadista de componentes eletrônicos e equipamentos de elefonia e comunicação 1652-4/00 - Comércio atacadista de componentes eletrônicos e equipamentos de elefonia e comunicação 1624-0/00 - Consultoria em tecnologia da informação 1629-1/00 - Suporte técnico, manutenção e outros serviços em tecnologia da informação 17490-1/04 - Atividades de intermediação e agenciamento de serviços e negócios em geral, exceto imobiliários 1733-1/00 - Aluguel de máquinas e equipamentos para escritórios veracidade da informação poderá ser verificada na seguinte página da Internet: https://www.serra.es.gov have de Verificação:	Lódigo e descrição da 1651-6/01 - Comé	atividade econômic ércio atacadista	a principal a de equipamento	os de informá	tica		
Alego descrição das atividades de internet 14652-4/00 - Comércio atacadista de componentes eletrônicos e equipamentos de telefonia e comunicação 20digo e descrição das atividades de licença 14652-4/00 - Comércio atacadista de componentes eletrônicos e equipamentos de telefonia e comunicação 20digo e descrição das atividades de serviço 20digo e descrição dos CNAEs 2651-6/01 - Comércio atacadista de equipamentos de informática 2651-6/02 - Comércio atacadista de suprimentos para informática 2652-4/00 - Comércio atacadista de componentes eletrônicos e equipamentos de elefonia e comunicação 2204-0/00 - Consultoria em tecnologia da informação 2209-1/00 - Suporte técnico, manutenção e outros serviços em tecnologia da informação 2733-1/00 - Aluguel de máquinas e equipamentos para escritórios veracidade da informação poderá ser verificada na seguinte página da Internet: https://www.serra.es.gov	ódigo e descrição da	a atividade de licenco	2				
Código e descrição das atividades de licença 4652-4/00 - Comércio atacadista de componentes eletrônicos e equipamentos de telefonia e comunicação Código e descrição das atividades de serviço Código e descrição das atividades de serviço Código e descrição das atividades de serviço Código e descrição dos CNAEs 4651-6/01 - Comércio atacadista de equipamentos de informática 4651-6/02 - Comércio atacadista de suprimentos para informática 4652-4/00 - Comércio atacadista de componentes eletrônicos e equipamentos de elefonia e comunicação 5204-0/00 - Consultoria em tecnologia da informação 5209-1/00 - Suporte técnico, manutenção e outros serviços em tecnologia da informação 7490-1/04 - Atividades de intermediação e agenciamento de serviços e negócios em geral, exceto imobiliários 7733-1/00 - Aluguel de máquinas e equipamentos para escritórios veracidade da informação poderá ser verificada na seguinte página da Internet: https://www.serra.es.gov have de Verificação:	4652-4/00 - Come	ércio atacadista	a de componente	es eletrônicos	e equipamentos de	e telefonia e c	omunicação
Código e descrição das atividades de serviço Código e descrição dos CNAEs 1651-6/01 - Comércio atacadista de equipamentos de informática 1651-6/02 - Comércio atacadista de suprimentos para informática 1652-4/00 - Comércio atacadista de componentes eletrônicos e equipamentos de elefonia e comunicação 3204-0/00 - Consultoria em tecnologia da informação 3209-1/00 - Suporte técnico, manutenção e outros serviços em tecnologia da informação 7490-1/04 - Atividades de intermediação e agenciamento de serviços e negócios em geral, exceto imobiliários 7733-1/00 - Aluguel de máquinas e equipamentos para escritórios veracidade da informação poderá ser verificada na seguinte página da Internet: https://www.serra.es.gov have de Verificação:	Código e descrição da 4652-4/00 - Com	as atividades de lice ércio atacadista	^{nça} a de componente	es eletrônicos	e equipamentos de	e telefonia e c	omunicação
Código e descrição dos CNAEs Data Encerramento 4651-6/01 - Comércio atacadista de equipamentos de informática 4651-6/02 - Comércio atacadista de suprimentos para informática 4651-6/02 - Comércio atacadista de componentes eletrônicos e equipamentos de elefonia e comunicação 6204-0/00 - Consultoria em tecnologia da informação 5209-1/00 - Suporte técnico, manutenção e outros serviços em tecnologia da informação 6209-1/00 - Suporte técnico, manutenção e agenciamento de serviços e negócios em geral, exceto imobiliários 7733-1/00 - Aluguel de máquinas e equipamentos para escritórios veracidade da informação poderá ser verificada na seguinte página da Internet: https://www.serra.es.gov have de Verificação: 100 - Suporte técnico, serviços en eleviços eleviços en eleviços en eleviços eleviços en eleviços eleviços en elevicitos eleviços en eleviços en eleviços en eleviços en eleviços en e	Código e descrição da	as atividades de serv	viço				
4651-6/01 - Comércio atacadista de equipamentos de informática 4651-6/02 - Comércio atacadista de suprimentos para informática 4652-4/00 - Comércio atacadista de componentes eletrônicos e equipamentos de elefonia e comunicação 5204-0/00 - Consultoria em tecnologia da informação 5209-1/00 - Suporte técnico, manutenção e outros serviços em tecnologia da informação 7490-1/04 - Atividades de intermediação e agenciamento de serviços e negócios em geral, exceto imobiliários 7733-1/00 - Aluguel de máquinas e equipamentos para escritórios	Código e descrição do	os CNAEs				Data	a Encerramento
4651-6/02 - Comércio atacadista de suprimentos para informática 4652-4/00 - Comércio atacadista de componentes eletrônicos e equipamentos de elefonia e comunicação 5204-0/00 - Consultoria em tecnologia da informação 5209-1/00 - Suporte técnico, manutenção e outros serviços em tecnologia da informação 7490-1/04 - Atividades de intermediação e agenciamento de serviços e negócios em geral, exceto imobiliários 7733-1/00 - Aluguel de máquinas e equipamentos para escritórios	1651-6/01 - Comé	ércio atacadista	a de equipamente	os de informá	tica		
4652-4/00 - Comércio atacadista de componentes eletrônicos e equipamentos de elefonia e comunicação 5204-0/00 - Consultoria em tecnologia da informação 5209-1/00 - Suporte técnico, manutenção e outros serviços em tecnologia da informação 7490-1/04 - Atividades de intermediação e agenciamento de serviços e negócios em geral, exceto imobiliários 7733-1/00 - Aluguel de máquinas e equipamentos para escritórios veracidade da informação poderá ser verificada na seguinte página da Internet: https://www.serra.es.gov have de Verificação:	1651-6/02 - Comé	ércio atacadista	a de suprimentos	s para informá	tica		
5204-0/00 - Consultoria em tecnologia da informação 5209-1/00 - Suporte técnico, manutenção e outros serviços em tecnologia da informação 7490-1/04 - Atividades de intermediação e agenciamento de serviços e negócios em geral, exceto imobiliários 7733-1/00 - Aluguel de máquinas e equipamentos para escritórios veracidade da informação poderá ser verificada na seguinte página da Internet: https://www.serra.es.gov	4652-4/00 - Comé elefonia e comu	ércio atacadista nicacão	a de componente	es eletrônicos	e equipamentos de		
6209-1/00 - Suporte técnico, manutenção e outros serviços em tecnologia da informação 7490-1/04 - Atividades de intermediação e agenciamento de serviços e negócios em geral, exceto imobiliários 7733-1/00 - Aluguel de máquinas e equipamentos para escritórios veracidade da informação poderá ser verificada na seguinte página da Internet: https://www.serra.es.gov	6204-0/00 - Cons	ultoria em tecn	ologia da inform	iação			
7490-1/04 - Atividades de intermediação e agenciamento de serviços e negócios em geral, exceto imobiliários 7733-1/00 - Aluguel de máquinas e equipamentos para escritórios veracidade da informação poderá ser verificada na seguinte página da Internet: https://www.serra.es.gov have de Verificação:	6209-1/00 - Supo	rte técnico, ma	nutenção e outro	os serviços er	n tecnologia da info	ormação	
7733-1/00 - Aluguel de máquinas e equipamentos para escritórios veracidade da informação poderá ser verificada na seguinte página da Internet: https://www.serra.es.gov have de Verificação:	7490-1/04 - Ativic geral, exceto imo	dades de intern obiliários	nediação e ageno	ciamento de s	erviços e negócios	em	
veracidade da informação poderá ser verificada na seguinte página da Internet: https://www.serra.es.gov have de Verificação:	7733-1/00 - Alugı	uel de máquina	s e equipamento	os para escritó	orios		
veracidade da informação poderá ser verificada na seguinte página da Internet: https://www.serra.es.go have de Verificação:						I	
have de Verificação:	veracidade da i	nformação pod	erá ser verificad	a na seguinte	página da Internet:	https://www.	serra.es.gov
	have de Verifi	cação:					

Inserido ao protocolo 22.951.206-4 por: Pedro Henrique Golin Linhares em: 31/03/2025 12:49. A autenticidade deste documento pode ser validada no endereço: https://www.eprotocolo.pr.gov.br/spiweb/validarDocumento com o código: 1d3e5a42275b5642ba40f346676ed94.




Ministério da Economia Secretaria de Governo Digital Departamento Nacional de Registro Empresarial e Integração				Nº ⊡	O PROTOCOLO (Uso da	a Junta Comercial)	PROTO	
Secre	etaria de Desei	nvolvimer	nto Econômico	o e Turismo				A Mov.
IIRE (da sede ou filia ede for em outra UF	l, quando a	Código da Jurídica	Natureza	Nº de Matrícula do A Auxiliar do Comércio	Agente o			TADO
4320378	9542	2	2062					
- REQUERIME	NTO							
	ILMO(A). S	SR.(A) I	PRESIDEN	ITE DA Junta C	Comercial, Ind	lustrial e Serviços do	Rio Grande do S	Sul
lome:	PERFIL COM	PUTACIC	<u>DNAL LTDA</u>					
	(da Empresa d	ou do Age	ente Auxiliar d	o Comércio)			Nº FCN/RE	MP
equer a V.Sª o de	ferimento do se	eguinte a	ito:					
INDE CODIGO	EVENTO	QTDE	DESCRIÇÃO	D DO ATO / EVEN	ТО		RSE2	300169452
002			ALTERACA	0				
	051	1	CONSOLID	ACAO DE CONTR)		
	026	1	ABERTURA	DE FILIAL EM OU	JTRA UF			
		E/ 1	ARROUPILHA Local 11 Maio 2023 Data	<u>.</u>	Representa Nome: Assina Telefor	tura:	/ Agente Auxiliar d	
- USO DA JUN		CIAL						
lome(s) Empresa	rial(ais) igual(a	is) ou ser	melhante(s):			JCOLEGIADA	1	
	(SIM			Process	o em Ordem
							À	decisão
							۸ () () () () () () () () () (decisão / Data
							À	decisão / Data
							À	decisão / Data
NÃO/	_/ Data	Res	ponsável		_// Data	Responsável	À (/	decisão / Data ponsável
NÃO/_	/ Data	Res	ponsável	NÃO	_// Data 2ª Exigência	Responsável	À i	decisão / Data ponsável 5º Exigência
NÃO/_ ECISÃO SINGUL Processo em	_/ Data _AR exigência. (Vid	Respland	ponsável	NÃO	_// Data 2ª Exigência	Responsável 3ª Exigência	À d /	decisão / Data ponsável 5º Exigência
NÃO/_ ECISÃO SINGUI Processo em Processo defe	/ Data AR exigência. (Vid erido. Publique	Resp de despac	ponsável cho em folha a uive-se.	. NÃO	_// Data 2ª Exigência	Responsável 3ª Exigência	À i /	decisão / Data ponsável 5ª Exigência
NÃO/_ ECISÃO SINGUL Processo em Processo defe Processo inde	/ Data _AR exigência. (Vid erido. Publique oferido. Publique	Resj de despac -se e arqu ue-se.	ponsável cho em folha a uive-se.	NÃO	// Data 2ª Exigência	Responsável 3ª Exigência	À i /	decisão / Data ponsável 5º Exigência
NÃO/_ ECISÃO SINGUI Processo em Processo defe Processo inde	/ Data _AR exigência. (Vid erido. Publique eferido. Publique	Resp de despac -se e arqu ue-se.	ponsável cho em folha a uive-se.	. NÃO	_// Data 2ª Exigência	Responsável 3ª Exigência	À (/	decisão Data ponsável 5ª Exigência
NÃO/_ ECISÃO SINGUI Processo em Processo defe Processo inde	/ Data _AR exigência. (Vid erido. Publique eferido. Publique	Resp de despac -se e arqu ue-se.	ponsável cho em folha a uive-se.	NÃO	// Data 2ª Exigência	Responsável 3ª Exigência	À (/	decisão Data ponsável 5º Exigência Responsável
NÃO/_ DECISÃO SINGUL Processo em Processo defe Processo inde DECISÃO COLEG Processo em	/ Data AR exigência. (Vid erido. Publique oferido. Publique IADA exigência. (Vid	Resp de despac -se e arqu ue-se. de despac	ponsável cho em folha a uive-se.	NÃO anexa)	_// Data 2ª Exigência	Responsável 3ª Exigência 3ª Exigência	À d /	decisão / Data ponsável 5ª Exigência Responsável 5ª Exigência
NÃO/_ DECISÃO SINGUI Processo em Processo inde DECISÃO COLEG Processo em Processo em Processo em	/ Data _AR exigência. (Vid erido. Publique eferido. Publique IADA exigência. (Vid erido. Publique	Resp de despac -se e arqu ue-se. de despac	ponsável cho em folha a uive-se. cho em folha a	. NÃO	 Data 2ª Exigência 2ª Exigência	Responsável 3ª Exigência 3ª Exigência 3ª Exigência	À d /	decisão Data ponsável \$^a Exigência Responsável 5^a Exigência
NÃO/_ ECISÃO SINGUI Processo em Processo defe Processo inde ECISÃO COLEG Processo em Processo defe Processo defe	/ Data 	Resp de despace -se e arqu Je-se. de despace -se e arqu Je-se.	ponsável cho em folha a uive-se. cho em folha a uive-se.	NÃO	_// Data 2ª Exigência 2ª Exigência	Responsável	À d /	decisão / Data ponsável 5ª Exigência Responsável 5ª Exigência Responsável 5ª Exigência
NÃO/_ ECISÃO SINGUI Processo em Processo defe Processo inde ECISÃO COLEG Processo em Processo defe Processo defe	/ Data _AR exigência. (Vid erido. Publique eferido. Publique exigência. (Vid erido. Publique eferido. Publique	Resp de despac -se e arqu ue-se. de despac -se e arqu ue-se.	ponsável cho em folha a uive-se.	. NÃO	 Data 2ª Exigência 2ª Exigência	Responsável 3ª Exigência 3ª Exigência 3ª Exigência	À i /_ 	decisão Data ponsável Responsável 5ª Exigência Responsável
NÃO/_ ECISÃO SINGUI Processo em Processo defe Processo inde ECISÃO COLEG Processo em Processo defe Processo inde	/ Data 	Res de despac -se e arqu ue-se. de despac -se e arqu ue-se.	ponsável cho em folha a uive-se. cho em folha a uive-se.	. NÃO	 Data 2ª Exigência 2ª Exigência 2ª Exigência	Responsável	À i /_ 	decisão / Data ponsável 5º Exigência Responsável 5º Exigência Componsável Sº Exigência Vogal
NÃO/_ ECISÃO SINGUI Processo em Processo inde Processo inde ECISÃO COLEG Processo em Processo defe Processo inde	/ Data 	Resp de despac -se e arqu ue-se. de despac -se e arqu ue-se.	ponsável cho em folha a uive-se.	NÃO		Responsável	À i /_ 	decisão / Data ponsável S ^a Exigência Responsável 5 ^a Exigência Vogal
NÃO/_ ECISÃO SINGUI Processo em Processo defe Processo inde ECISÃO COLEG Processo em Processo defe Processo inde	/ Data 	Res de despac -se e arqu ue-se. de despac -se e arqu ue-se.	ponsável cho em folha a uive-se.	anexa)		Responsável 3ª Exigência 3ª Exigência 3ª Exigência Vogal a Turma	À i /_ 	decisão / Data ponsável 5º Exigência Responsável 5º Exigência Compondante Vogal
NÃO/_ DECISÃO SINGUI Processo em Processo defe Processo inde DECISÃO COLEG Processo em Processo defe Processo defe OFOCES	/ Data 	Resp de despace -se e arqu ue-se. de despace -se e arqu ue-se.	ponsável cho em folha a uive-se.	NÃO anexa)	 Data 2ª Exigência 2ª Exigência 2ª Exigência 2ª Exigência Vogal Presidente da	Responsável Responsável	À i /_ 	decisão / Data ponsável 5º Exigência Responsável 5º Exigência Vogal Vogal
NÃO/_ ECISÃO SINGUI Processo em Processo inde ECISÃO COLEG Processo em Processo inde Necesso inde BECISÃO COLEG	/ Data 	Resp de despac -se e arqu ue-se. de despac -se e arqu ue-se.	ponsável cho em folha a uive-se.	anexa)	 Data 2ª Exigência 2ª Exigência 2ª Exigência Vogal Presidente da	Responsável Responsável	À i /_ 	decisão / Data ponsável 5º Exigência Responsável 5º Exigência Vogal Vogal
NÃO/_ ECISÃO SINGUI Processo em Processo defe Processo inde ECISÃO COLEG Processo em Processo defe Processo inde Processo inde	/ Data 	Resp de despace -se e arqu ue-se. de despace -se e arqu ue-se.	ponsável cho em folha a uive-se.		 Data 2ª Exigência 2ª Exigência 2ª Exigência Vogal Presidente da	Responsável Responsável 3ª Exigência 3ª Exigência 3ª Exigência Vogal aTurma	À i /_ 	decisão / Data ponsável 5º Exigência Responsável 5º Exigência Vogal Vogal

Certifico registro sob o nº 8964803 em 31/05/2023 da Empresa PERFIL COMPUTACIONAL LTDA, CNPJ 02543216000129 e protocolo 231521995 - 30/05/2023. Autenticação: F2F87A17B71DA2A5B81B99D73685613DD765BF. José Tadeu Jacoby - Secretário-Geral. Para validar este documento, acesse http://jucisrs.rs.gov.br/validacao e informe nº do protocolo 23/152.199-5 e o código de segurança 7LAQ Esta cópia foi autenticada digitalmente e assinada em 02/06/2023 por José Tadeu Jacoby Secretário-Geral.

pág. 1/11

Inserido ao protocolo **22.951.206-4** por: **Pedro Henrique Golin Linhares** em: 31/03/2025 12:49. A autenticidade deste documento pode ser validada no endereço: https://www.eprotocolo.pr.gov.br/spiweb/validarDocumento com o código: **1d3e5a42275b5642ba40f346676ed94**.

5000



JUNTA COMERCIAL, INDUSTRIAL E SERVIÇOS DO RIO GRANDE DO SUL

Registro Digital



Capa de Processo

Identificação do Pro	cesso	
Número do Protocolo	Número do Processo Módulo Integrador	Data
23/152.199-5	RSE2300169452	11/05/2023

Identificação do(s	s) Assinante(s)	
CPF	Nome	Data Assinatura
805.127.710-49	IGOR SIDNEI REOLON	12/05/2023
		1317

Assinado utilizando o(s) seguinte(s) selo(s) do govbr 🤊

Selo Prata - Biometria Facial, Selo Prata - Banco do Brasil - Internet Banking, Selo Ouro - Certificado Digital, Selo Prata - Cadastro via Internet Banking



Junta Comercial, Industrial e Serviços do Rio Grande do Sul Certifico registro sob o nº 8964803 em 31/05/2023 da Empre

Certifico registro sob o nº 8964803 em 31/05/2023 da Empresa PERFIL COMPUTACIONAL LTDA, CNPJ 02543216000129 e protocolo 231521995 - 30/05/2023. Autenticação: F2F87A17B71DA2A5B81B99D73685613DD765BF. José Tadeu Jacoby - Secretário-Geral. Para validar este documento, acesse http://jucisrs.rs.gov.br/validacao e informe nº do protocolo 23/152.199-5 e o código de segurança 7LAQ Esta cópia foi autenticada digitalmente e assinada em 02/06/2023 por José Tadeu Jacoby Secretário-Geral.

pág. 2/11



PERFIL COMPUTACIONAL LTDA

Rua Barão do Rio Branco, nº 459, Sala 20

Bairro Centro – CEP 95170-404

FARROUPILHA – RIO GRANDE DO SUL

CNPJ: 02.543.216/0001-29 NIRE: 43.203.789.542

INSTRUMENTO PARTICULAR DE VIGÉSIMA QUARTA ALTERAÇÃO E CONSOLIDAÇÃO DE CONTRATO SOCIAL

Pelo presente instrumento particular de alteração de contrato social, e na melhor forma de direito permitida, os abaixo assinados:

IGOR SIDNEI REOLON, brasileiro, casado sob o regime de comunhão parcial de bens, nascido em 18.09.1980, empresário, residente e domiciliado na Rua Barão do Rio Branco, nº 459, Apto 401, Bairro Centro, em Farroupilha, RS, CEP 95170-404, inscrito no CPF sob nº 805.127.710-49, portador da Carteira de Identidade RG nº 1067955946, expedida pela SJTC/RS, em 15.03.1996;

e,

FÁBIO VITOR REOLON, brasileiro, solteiro, maior, nascido em 24.04.1985, empresário, residente e domiciliado na Rua Barão do Rio Branco, nº 459, Apto 202, Bairro Centro, em Farroupilha, RS, CEP 95170-404 inscrito no CPF sob nº 010.120.500-76, e portador da Carteira de Identidade RG nº 6090528875, expedida pela SSP/RS.

... sócios componentes da sociedade que gira sob a denominação social **PERFIL COMPUTACIONAL LTDA**, com sede na Rua Barão do Rio Branco, nº 459, Sala 20, Bairro Centro, em Farroupilha, RS, CEP 95170-404, inscrita no CNPJ sob nº **02.543.216/0001-29**, constituída conforme contrato social arquivado na MM. Junta Comercial, Industrial e Serviços do Rio Grande do Sul sob NIRE nº **43.203.789.542**, em sessão de 19 de maio de 1998, e alterações posteriores, resolvem de comum acordo alterar e posteriormente consolidar o seu instrumento particular de contrato social, fazendo-o neste ato da seguinte forma:

DA ABERTURA DE FILIAL

É aberta, neste ato, a nova filial com o mesmo ramo de atividade da matriz, no respectivo logradouro:

Filial Goiânia Avenida 136, nº 761, 14º andar, Sala B142, Caixa Postal W1, Quadra F-44, Lote 2E, Edifício Nasa Business Style, Bairro Setor Sul, em Goiânia, GO, CEP 74.093-250.

CONSOLIDAÇÃO DO CONTRATO SOCIAL

Pelo presente instrumento particular de consolidação contratual, e na melhor forma de direito permitida, os abaixo assinados:

IGOR SIDNEI REOLON, brasileiro, casado sob o regime de comunhão parcial de bens, nascido em 18.09.1980, empresário, residente e domiciliado na Rua Barão do Rio Branco, nº 459, Apto 401, Bairro Centro, em Farroupilha, RS, CEP 95170-404, inscrito no CPF sob nº 805.127.710-49, portador da Carteira de Identidade RG nº 1067955946, expedida pela SJTC/RS, em 15.03.1996;

e,

Junta Comercial, Industrial e Serviços do Rio Grande do Sul

Certifico registro sob o nº 8964803 em 31/05/2023 da Empresa PERFIL COMPUTACIONAL LTDA, CNPJ 02543216000129 e protocolo 231521995 - 30/05/2023. Autenticação: F2F87A17B71DA2A5B81B99D73685613DD765BF. José Tadeu Jacoby - Secretário-Geral. Para validar este documento, acesse http://jucisrs.rs.gov.br/validacao e informe nº do protocolo 23/152.199-5 e o código de segurança 7LAQ Esta cópia foi autenticada digitalmente e assinada em 02/06/2023 por José Tadeu Jacoby Secretário-Geral.

pág. 3/11



FÁBIO VITOR REOLON, brasileiro, solteiro, maior, nascido em 24.04.1985, empresário, residente e domiciliado na Rua Barão do Rio Branco, nº 459, Apto 202, Bairro Centro, em Farroupilha, RS, CEP 95170-404, inscrito no CPF sob nº 010.120.500-76, e portador da Carteira de Identidade RG nº 6090528875, expedida pela SSP/RS.

... resolvem de comum acordo consolidar o seu instrumento particular de alteração de contrato social, fazendo-o neste ato e da seguinte forma:

<u>Primeira</u>

A sociedade limitada gira sob a denominação social de **PERFIL COMPUTACIONAL LTDA**, da qual podem fazer uso os sócios adiante designados como administradores.

<u>Segunda</u>

A sede da sociedade está localizada na Rua Barão do Rio Branco, nº 459, Sala 20, Bairro Centro, em Farroupilha, RS, CEP 95170-404.

A sociedade possui a seguinte filial, com o objeto social de: a) comércio atacadista de equipamentos de informática (4651-6/01), b) suporte técnico, manutenção e outros serviços em tecnologia da informação, que compreende: os serviços de instalação de equipamentos de informática (hardware), os serviços de instalação de programas de computador (software), o suporte técnico em tecnologia da informação (6209-1/00), c) aluguel de máquinas e equipamentos para escritórios, que compreende: o aluguel de equipamentos para processamento de dados, o aluguel de equipamentos e periféricos de informática (7733-1/00), d) comércio atacadista de suprimentos para informática (4651-6/02); e) consultoria em tecnologia da informação (6204-0/00), f) comércio atacadista de componentes eletrônicos e equipamentos de telefonia e comunicação (4652-4/00).

Filial Barueri Alameda Rio Negro, n° 1030, Escritório 206, Sala 03, Condomínio Stadium, Bairro Alphaville Centro Industrial e empresarial - Barueri - São Paulo - CEP 06454-000, devidamente inscrita no CNPJ sob n° 02.543.216/0010-10 e NIRE n° 35.905.941.194.

E as filiais abaixo, com o objeto social de: a) comércio atacadista de equipamentos de informática (4651-6/01), b) suporte técnico, manutenção e outros serviços em tecnologia da informação, que compreende: os serviços de instalação de equipamentos de informática (hardware), os serviços de instalação de programas de computador (software), o suporte técnico em tecnologia da informação (6209-1/00), c) aluguel de máquinas e equipamentos para escritórios, que compreende: o aluguel de equipamentos para processamento de dados, o aluguel de equipamentos e periféricos de informática (7733-1/00), d) comércio atacadista de suprimentos para informática (4651-6/02); e) consultoria em tecnologia da informação (6204-0/00), f) comércio atacadista de componentes eletrônicos e equipamentos de telefonia e comunicação (4652-4/00); g) intermediação e agenciamento de serviços e negócios, exceto imobiliários (7490-1/04).

Filial Serra	Avenida Acesso Rodoviário, S/N, Quadra 06, L-M01, Quadra 01, L-M18 a
	M23, Sala 107, Terminal Intermodal da Serra, Serra, ES, CEP 29.161-376,
	devidamente inscrita no CNPJ sob o nº 02.543.216/0011-09 e NIRE nº
	32.900.792.252.

- Filial Brasília SCS Quadra 06, Bloco A, nº 130, 6º andar, parte H2, Edifício Ermes, em Brasília, DF, CEP 70306-901, devidamente inscrita no CNPJ sob n° 02.543.216/0007-14 e NIRE n°53.900.356.328.
- **Filial Curitiba** Rua Antônio Lacerda Braga, nº 960, Sala J, Cidade Industrial Curitiba, em Curitiba, PR, CEP 81170-240, devidamente inscrita no CNPJ sob n° 02.543.216/0008-03 e NIRE n° 41.901.686.313.
- Filial Florianópolis Av. Rio Branco, nº 404, Torre II, Sala 908-A, Bairro Centro, em Florianópolis, SC, CEP 88015-200, devidamente inscrita no CNPJ sob n° 02.543.216/0006-33 e NIRE n° 42.901.114.281.

Junta Comercial, Industrial e Serviços do Rio Grande do Sul

Certifico registro sob o nº 8964803 em 31/05/2023 da Empresa PERFIL COMPUTACIONAL LTDA, CNPJ 02543216000129 e protocolo 231521995 - 30/05/2023. Autenticação: F2F87A17B71DA2A5B81B99D73685613DD765BF. José Tadeu Jacoby - Secretário-Geral. Para validar este documento, acesse http://jucisrs.rs.gov.br/validacao e informe nº do protocolo 23/152.199-5 e o código de segurança 7LAQ Esta cópia foi autenticada digitalmente e assinada em 02/06/2023 por José Tadeu Jacoby Secretário-Geral.

pág. 4/11



Filial Goiânia

Avenida 136, nº 761, 14º andar, Sala B142, Caixa Postal W1, Quadra F-44, Lote 2E, Edifício Nasa Business Style, Bairro Setor Sul, em Goiânia, GO, CEP 74.093-250.



<u>Terceira</u>

O objeto social da sociedade matriz é: a) comércio atacadista de equipamentos de informática (4651-6/01), b) suporte técnico, manutenção e outros serviços em tecnologia da informação, que compreende: os serviços de instalação de equipamentos de informática (hardware), os serviços de instalação de programas de computador (software), o suporte técnico em tecnologia da informação (6209-1/00), c) aluguel de máquinas e equipamentos para escritórios, que compreende: o aluguel de equipamentos para processamento de dados, o aluguel de equipamentos e periféricos de informática (7733-1/00), d) comércio atacadista de suprimentos para informática (4651-6/02); e) consultoria em tecnologia da informação (6204-0/00), f) comércio atacadista de componentes eletrônicos e equipamentos de telefonia e comunicação (4652-4/00); g) intermediação e agenciamento de serviços e negócios, exceto imobiliários (7490-1/04).

Quarta

A sociedade é administrada pelo sócio **IGOR SIDNEI REOLON**, individualmente, ao qual cabe a representação da sociedade em juízo e/ou fora dele, podendo praticar todos os atos inerentes e relativos ao objeto da sociedade, com os mais amplos e irrestritos poderes para a administração, assumindo aquelas funções, independentemente de caução ou qualquer outra formalidade, ficando, entretanto, expressamente proibido o uso da mesma em fianças, abonos, endossos ou quaisquer outras transações alheias ou estranhas às finalidades da sociedade.

- § 1º Para admitir e demitir empregados, representar a sociedade perante órgãos públicos, em licitações e concorrências, abrir e fechar filiais em qualquer parte do território nacional, vender, alienar ou onerar qualquer bem do ativo fixo da sociedade, contratar empréstimos ou financiamentos bancários, nomear procuradores ad judicia e ad negotia, será sempre necessária, indispensável e suficiente, a assinatura do sócio administrador, individualmente.
- § 2º Para assinar contratos de certificação digital, não podendo o administrador fazê-lo, fica autorizada a representação por meio de procuração por instrumento público.
- § 3º Declara o administrador, sob as penas da lei, que não está impedido por lei especial de exercer a administração da sociedade, nem condenado ou sob efeitos de condenação, a pena que vede, ainda que temporariamente, o acesso a cargos públicos, ou por crime falimentar, de prevaricação, peita ou suborno, concussão, peculato, ou contra a economia popular, contra o sistema financeiro nacional, contra as normas de defesa da concorrência, contra as relações de consumo, a fé pública ou a propriedade.
- § 4º A sociedade poderá ser administrada por não sócio, desde que sua designação seja aprovada pela unanimidade dos sócios.

Quinta

A sociedade poderá a qualquer tempo, abrir, fechar, alterar o objeto social, endereços, dentre outras alterações pertinentes as filiais, mediante Ata de Reunião de Sócios.

<u>Sexta</u>

Os sócios que exercerem atividades na sociedade terão direito a uma remuneração mensal a título de "pró-labore", previamente combinado entre todos os participantes do capital social.

<u>Sétima</u>

O capital social da sociedade é da importância de R\$ 15.000.000,00 (quinze milhões) dividido em 15.000.000 (quinze milhões) quotas no valor de R\$ R\$ 1,00 (um real) cada, totalmente subscrito e integralizado em moeda corrente nacional, fica a seguir melhor distribuído entre os sócios, a saber:

Junta Comercial, Industrial e Serviços do Rio Grande do Sul

Certifico registro sob o nº 8964803 em 31/05/2023 da Empresa PERFIL COMPUTACIONAL LTDA, CNPJ 02543216000129 e protocolo 231521995 - 30/05/2023. Autenticação: F2F87A17B71DA2A5B81B99D73685613DD765BF. José Tadeu Jacoby - Secretário-Geral. Para validar este documento, acesse http://jucisrs.rs.gov.br/validacao e informe nº do protocolo 23/152.199-5 e o código de segurança 7LAQ Esta cópia foi autenticada digitalmente e assinada em 02/06/2023 por José Tadeu Jacoby Secretário-Geral.

pág. 5/11



SÓCIOS	QUOTAS		VALOR	%
IGOR SIDNEI REOLON	13.650.000	R\$	13.650.000,00	91,00%
FÁBIO VITOR REOLON	1.350.000	R\$	1.350.000,00	9,00%
TOTAL	15.000.000	R\$	15.000.000,00	100%



Parágrafo único.

A responsabilidade de cada sócio é restrita ao valor de suas quotas, mas todos respondem solidariamente pela integralização do capital social.

<u>Oitava</u>

O prazo de duração da sociedade é por tempo indeterminado e suas atividades tiveram início em 07 de maio de 1998.

<u>Nona</u>

As deliberações dos sócios serão tomadas em reuniões, as quais se tornam dispensáveis quando todos os sócios decidirem, por escrito, sobre a matéria que seria objeto delas.

- § 1º Dispensam-se as formalidades de convocação previstas no § 3º do Art. 1.152 da Lei 10.406 de 10 de janeiro de 2002, quando todos os sócios comparecerem ou se declararem, por escrito, cientes do local, data, hora e ordem do dia.
- § 2º O sócio pode ser representado na reunião por outro sócio ou por advogado mediante apresentação de instrumento de mandato. Na hipótese de falecimento do sócio, a sua representação na reunião de sócios será exercida pelo inventariante nomeado.

<u>Décima</u>

Anualmente proceder-se-á, ao término do exercício social, que ocorrerá em 31 (trinta e um) de dezembro, a elaboração do inventário, do balanço patrimonial e do balanço de resultado econômico, sendo realizada nos quatro meses seguintes ao término do exercício social uma reunião dos sócios para deliberar sobre os mesmos e sobre as contas dos administradores.

- § 1º Até 30 (trinta) dias antes da data marcada para a reunião, os balanços sociais devem ser postos à disposição dos sócios que não exerçam a administração.
- § 2º Dos trabalhos e deliberações será lavrada ata, assinada pelos sócios participantes da reunião; cópia desta ata, autenticada pelos administradores, será, nos 20 (vinte) dias subseqüentes à reunião apresentada à Junta Comercial competente para averbação e arquivamento.
- § 3º A sociedade, a critério dos sócios, poderá realizar o levantamento de balanços intermediários, os quais deverão atender as formalidades legais, em períodos menores do que o previsto no "caput" desta cláusula, para o fim específico de distribuição de lucros.

Décima Primeira

Os lucros verificados serão partilhados proporcionalmente ao capital social, se não for deliberado e acordado diversamente em reunião de sócios, podendo, antes, serem criados os fundos necessários e legais para a sociedade. Em caso de perdas (prejuízos), estas serão mantidas pendentes para compensação em exercícios posteriores ou partilhadas proporcionalmente ao capital social.

Décima Segunda

Em caso de falecimento de qualquer um dos sócios, a sociedade não se dissolverá, podendo continuar com o remanescente e os herdeiros legais do sócio falecido, caso seja do interesse dos mesmos. Não havendo interesse por parte dos herdeiros do sócio pré-morto em continuar na sociedade, estes serão pagos de seus haveres, mediante o levantamento de um Balanço Geral Especial, na data do evento morte, sendo o *quantum* apurado correspondente à participação societária do "de cujus" pago a quem de direito em até 36 (trinta e seis) parcelas consecutivas mensalmente, corrigidas pelo IGP-M, ou qualquer outro índice criado pela legislação competente que o substitua, acrescidas de juros de 1% (um por cento) ao mês. A primeira parcela

Junta Comercial, Industrial e Serviços do Rio Grande do Sul

Certifico registro sob o nº 8964803 em 31/05/2023 da Empresa PERFIL COMPUTACIONAL LTDA, CNPJ 02543216000129 e protocolo 231521995 - 30/05/2023. Autenticação: F2F87A17B71DA2A5B81B99D73685613DD765BF. José Tadeu Jacoby - Secretário-Geral. Para validar este documento, acesse http://jucisrs.rs.gov.br/validacao e informe nº do protocolo 23/152.199-5 e o código de segurança 7LAQ Esta cópia foi autenticada digitalmente e assinada em 02/06/2023 por José Tadeu Jacoby Secretário-Geral.

pág. 6/11





terá vencimento em 30 (trinta) dias contados da data do Balanço Geral Especial e as demais no mesmo dia dos meses subsequentes.

Décima Terceira

Na hipótese de qualquer um dos sócios desejar ceder ou transferir parte ou a totalidade de suas quotas sociais, deverá primeiramente oferecê-las por escrito à sociedade e aos demais sócios, informando o preço e as condições para a venda, os quais terão o prazo máximo de 30 (trinta) dias, comum, para exercerem o direito de preferência que esta cláusula lhes assegura. Decorrido este prazo e não havendo interesse manifestado pela sociedade e nem pelos demais sócios na aquisição das quotas à venda, estas poderão ser oferecidas a terceiros, estranhos à sociedade, em nenhuma hipótese em condições mais favoráveis que as já apresentadas, e desde que estes, os terceiros, reúnam as condições jurídicas necessárias e sejam aceitos na sociedade, mediante a anuência expressa de todos os demais sócios.

Parágrafo único. Na hipótese de retirada do sócio da sociedade, sua quota será liquidada e paga nos mesmos termos da cláusula que trata do falecimento de sócio.

Décima Quarta

A maioria representativa de mais da metade do capital social poderá excluir por justa causa, mediante alteração do contrato social, o sócio que estiver pondo em risco a continuidade da empresa em virtude de atos de inegável gravidade.

- § 1º A exclusão de que trata esta cláusula somente poderá ser determinada em reunião especialmente convocada para este fim, ciente o acusado em tempo hábil para permitir seu comparecimento e o exercício do direito de defesa.
- § 2º O valor da quota do sócio excluído, considerada pelo montante efetivamente realizado, será apurada da mesma forma prevista na cláusula deste contrato que trata do falecimento de sócio, tomando-se como data base para o levantamento do Balanço Geral Especial a data da reunião que deliberou a exclusão. O pagamento dos haveres ao sócio excluído também obedecerá aos mesmos prazos e procedimentos da mencionada cláusula.

Décima Quinta

Fica eleito de comum acordo o Foro da Comarca de Farroupilha, RS, como competente para dirimir qualquer questão porventura decorrente deste instrumento, com renúncia expressa de qualquer outro por mais especial que possa ser.

Décima Sexta

Os casos omissos neste instrumento serão regulados pela Lei nº 10.406 de 10 de janeiro de 2002, supletivamente pela Lei 6.404 de 15 de dezembro de 1976 e pelos demais dispositivos legais incidentes sobre a matéria.

Décima Sétima

As cláusulas do contrato primitivo tornam-se sem efeito, passando, a partir desta data, a vigorar somente as cláusulas do presente instrumento particular de vigésima quarta alteração e consolidação de contrato social.

E, por assim estarem justos e contratados, obrigam-se por si e seus herdeiros a cumprirem fielmente o presente instrumento particular de contrato social, que assinam abaixo, seguindo-se posteriormente as demais exigências legais.

Farroupilha, RS, 05 de maio de 2023.

IGOR SIDNEI REOLON

FABIO VITOR REOLON

Junta Comercial, Industrial e Serviços do Rio Grande do Sul

Certifico registro sob o nº 8964803 em 31/05/2023 da Empresa PERFIL COMPUTACIONAL LTDA, CNPJ 02543216000129 e protocolo 231521995 - 30/05/2023. Autenticação: F2F87A17B71DA2A5B81B99D73685613DD765BF. José Tadeu Jacoby - Secretário-Geral. Para validar este documento, acesse http://jucisrs.rs.gov.br/validacao e informe nº do protocolo 23/152.199-5 e o código de segurança 7LAQ Esta cópia foi autenticada digitalmente e assinada em 02/06/2023 por José Tadeu Jacoby Secretário-Geral.

pág. 7/11





JUNTA COMERCIAL, INDUSTRIAL E SERVIÇOS DO RIO GRANDE DO SUL

Registro Digital

Fis. 279 Mov. 31

Documento Principal

Identificação do Pro	ocesso		
Número do Protocolo	Número do Processo Módulo Integrador		
23/152.199-5	RSE2300169452	11/05/2023	
Identificação do(s)	Assinante(s)		
CPF	Nome	1	Data Assinatura
010.120.500-76	FABIO VITOR REOLON	76	30/05/2023
Assinado utilizando o(s	s) seguinte(s) selo(s) do govbr 🔊 🕅	1 1.3	
Selo Ouro - Certificado	Digital, Selo Ouro - Biometria TSE, Selo Prata	- Biometria Facial	×
805.127.710-49	IGOR SIDNEI REOLON		12/05/2023
Assinado utilizando o(s	s) seguinte(s) selo(s) do govbr 🔊 🖽	(a).	
Selo Prata - Biometria via Internet Banking, S	Facial, Selo Prata - Banco do Brasil - Internet E elo Ouro - Certificado Digital	3anking, Selo Prata - 0	Cadastro



Junta Comercial, Industrial e Serviços do Rio Grande do Sul

Certifico registro sob o nº 8964803 em 31/05/2023 da Empresa PERFIL COMPUTACIONAL LTDA, CNPJ 02543216000129 e protocolo 231521995 - 30/05/2023. Autenticação: F2F87A17B71DA2A5B81B99D73685613DD765BF. José Tadeu Jacoby - Secretário-Geral. Para validar este documento, acesse http://jucisrs.rs.gov.br/validacao e informe nº do protocolo 23/152.199-5 e o código de segurança 7LAQ Esta cópia foi autenticada digitalmente e assinada em 02/06/2023 por José Tadeu Jacoby Secretário-Geral.

pág. 8/11



Ministério da Economia Secretaria de Governo Digital Departamento Nacional de Registro Empresarial e Integração Secretaria de Desenvolvimento Econômico e Turismo Junta Comercial, Industrial e Serviços do Rio Grande do Sul



Relatório de Filiais Abertas

Informamos que, do processo 23/152.199-5 arquivado nesta Junta Comercial sob o número 8964803 em 31/05/2023 da empresa 4320378954-2 PERFIL COMPUTACIONAL LTDA, consta a abertura da(s) seguinte(s) filial(ais):

NIRE	ENDEREÇO
xxxxxxx	AVENIDA 136 761 SALA B142 CXPST W1 QUADRAF-44 LOTE 2E EDIF NASA BUSINESS STYLE - BAIRRO SET SUL CEP 74093-250 - GOIANIA/GO

31 de mai de 2023

8

Junta Comercial, Industrial e Serviços do Rio Grande do Sul

Certifico registro sob o nº 8964803 em 31/05/2023 da Empresa PERFIL COMPUTACIONAL LTDA, CNPJ 02543216000129 e protocolo 231521995 - 30/05/2023. Autenticação: F2F87A17B71DA2A5B81B99D73685613DD765BF. José Tadeu Jacoby - Secretário-Geral. Para validar este documento, acesse http://jucisrs.rs.gov.br/validacao e informe nº do protocolo 23/152.199-5 e o código de segurança 7LAQ Esta cópia foi autenticada digitalmente e assinada em 02/06/2023 por José Tadeu Jacoby Secretário-Geral.

рад. 9/11



Sistema Nacional de Registro de Empresas Mercantil - SINREM Governo do Estado do Rio Grande Do Sul Secretaria de Desenvolvimento Econômico e Turismo Junta Comercial, Industrial e Serviços do Rio Grande do Sul



TERMO DE AUTENTICAÇÃO - REGISTRO DIGITAL

Certifico que o ato, assinado digitalmente, da empresa PERFIL COMPUTACIONAL LTDA, de CNPJ 02.543.216/0001-29 e protocolado sob o número 23/152.199-5 em 30/05/2023, encontra-se registrado na Junta Comercial sob o número 8964803, em 31/05/2023. O ato foi deferido eletronicamente pelo examinador Mario Ederich Filho.

Certifica o registro, o Secretário-Geral, José Tadeu Jacoby. Para sua validação, deverá ser acessado o sitio eletrônico do Portal de Serviços / Validar Documentos (https://portalservicos.jucisrs.rs.gov.br/Portal/pages/imagemProcesso/ viaUnica.jsf) e informar o número de protocolo e chave de segurança. Capa de Processo

Assinante(s)			
CPF	Nome	Data Assinatura	
805.127.710-49	IGOR SIDNEI REOLON	12/05/2023	
		P	

Assinado utilizando o(s) seguinte(s) selo(s) do govbr @

Selo Prata - Biometria Facial, Selo Prata - Banco do Brasil - Internet Banking, Selo Ouro - Certificado Digital, Selo Prata - Cadastro via Internet Banking

Documento Principal

Assinante(s)			
CPF	Nome	Data Assinatura	
805.127.710-49	IGOR SIDNEI REOLON	12/05/2023	
Assinado utilizando o	o(s) seguinte(s) selo(s) do govbr 🖉 🛄		

Selo Prata - Biometria Facial, Selo Prata - Banco do Brasil - Internet Banking, Selo Prata - Cadastro via Internet Banking, Selo Ouro - Certificado Digital 30/05/2023

010.120.500-76 FABIO VITOR REOLON

Assinado utilizando o(s) seguinte(s) selo(s) do govbr 🖉 🖽

Selo Ouro - Certificado Digital, Selo Ouro - Biometria TSE, Selo Prata - Biometria Facial

Data de início dos efeitos do registro (art. 36, Lei 8.934/1994): 05/05/2023



Documento assinado eletronicamente por Mario Ederich Filho, Servidor(a) Público(a), em 31/05/2023, às 17:38.



A autencidade desse documento pode ser conferida no portal de serviços da jucisrs informando o número do protocolo 23/152.199-5.

Junta Comercial, Industrial e Serviços do Rio Grande do Sul

Certifico registro sob o nº 8964803 em 31/05/2023 da Empresa PERFIL COMPUTACIONAL LTDA, CNPJ 02543216000129 e protocolo 231521995 -30/05/2023. Autenticação: F2F87A17B71DA2A5B81B99D73685613DD765BF. José Tadeu Jacoby - Secretário-Geral. Para validar este documento, acesse http://jucisrs.rs.gov.br/validacao e informe nº do protocolo 23/152.199-5 e o código de segurança 7LAQ Esta cópia foi autenticada digitalmente e assinada em 02/06/2023 por José Tadeu Jacoby Secretário-Geral.

pág. 10/11



JUNTA COMERCIAL, INDUSTRIAL E SERVIÇOS DO RIO GRANDE DO SUL Registro Digital



O ato foi assinado digitalmente por :

Identificação do(s)) Assinante(s)
CPF	Nome
054.744.500-87	JOSE TADEU JACOBY
	* VIII *

Porto Alegre. quarta-feira, 31 de maio de 2023

Junta Comercial, Industrial e Serviços do Rio Grande do Sul

Certifico registro sob o nº 8964803 em 31/05/2023 da Empresa PERFIL COMPUTACIONAL LTDA, CNPJ 02543216000129 e protocolo 231521995 - 30/05/2023. Autenticação: F2F87A17B71DA2A5B81B99D73685613DD765BF. José Tadeu Jacoby - Secretário-Geral. Para validar este documento, acesse http://jucisrs.rs.gov.br/validacao e informe nº do protocolo 23/152.199-5 e o código de segurança 7LAQ Esta cópia foi autenticada digitalmente e assinada em 02/06/2023 por José Tadeu Jacoby Secretário-Geral.

pág. 11/11

Voltar

Imprimir





Certificado de Regularidade do FGTS - CRF

Inscrição: Razão Social: Endereço:

02.543.216/0011-09 PERFIL COMPUTACIONAL LTDA AV ACESSO RODOVIARIO S/N / TERMINAL INTERMODAL / SERRA / ES / 29161-376

A Caixa Econômica Federal, no uso da atribuição que lhe confere o Art. 7, da Lei 8.036, de 11 de maio de 1990, certifica que, nesta data, a empresa acima identificada encontra-se em situação regular perante o Fundo de Garantia do Tempo de Servico - FGTS.

O presente Certificado não servirá de prova contra cobrança de quaisquer débitos referentes a contribuições e/ou encargos devidos, decorrentes das obrigações com o FGTS.

Validade:22/03/2025 a 20/04/2025

Certificação Número: 2025032200420852117414

Informação obtida em 24/03/2025 08:33:53

A utilização deste Certificado para os fins previstos em Lei esta condicionada a verificação de autenticidade no site da Caixa: **www.caixa.gov.br**



TABELIONATO KUNZLER

Daicir José Kunzler - Tabelião de Notas Rua Cel. Pena De Moraes, nº 661Z - Centro - Farroupilha/RS - 95.170-488 54 32611281



Nº 10.160 - PROCURAÇÃO, bastante que faz nestas notas, como outorgante, a empresa PERFIL COMPUTACIONAL LTDA, na forma a seguir. SAIBAM, todos quantos esta pública procuração virem que, aos trinta (30) dias do mês de junho, do ano de dois mil e vinte e três (2023), nesta cidade de Farroupilha, Estado do Rio Grande do Sul, neste Tabelionato, compareceu como outorgante, a empresa PERFIL COMPUTACIONAL LTDA., sociedade inscrita no CNPJ sob número 02.543.216/0001-29, com sede na Rua Barão do Rio Branco, nº 459, Sala 20, Centro, nesta cidade de Farroupilha/RS., com seu Contrato Social arquivado na Junta Comercial do Estado do Rio Grande do Sul sob NIRE número 43.203.789.542, em data de 19 de maio de 1998, e demais alterações posteriores, sendo a última, a 24ª Alteração e Consolidação do Contrato Social firmada em 05 de maio de 2023, devidamente registrada na mesma Junta Comercial sob número 8964803, em data de 31 de maio de 2023, neste ato, representada por seu sócio administrador, IGOR SIDNEI REOLON, brasileiro, empresário, portador da carteira nacional de habilitação sob número 00405237308, expedida pelo DETRAN/RS em 16/07/2020, inscrito no CPF sob número 805.127.710-49, filho de Vitorio Reolon e de Inês Balbinot Reolon, com endereço eletrônico: igor@perfil.inf.br, casado pelo regime da Comunhão Parcial de Bens posterior a Lei nº 6.515 de 26.12.77, com GISELI CRISTINA BELLAVER REOLON, residente e domiciliado na Rua Barão do Rio Branco, nº 459, apartamento 301, Centro, nesta cidade de Farroupilha/RS., nos termos da 24ª Alteração e Consolidação do Contrato Social, devidamente registrado neste Tabelionato sob nº 1674, às folhas 189 à 191, do Livro número 049 de Registro de Procurações, Autorizações Judiciais e Documentos de Representações Legais, em data de 14/06/2023, o qual declara que não houve outra alteração contratual da referida empresa até a presente data que possa descaracterizar sua representação jurídica. O comparecente, identificado documentalmente como o próprio por mim, DAICIR JOSÉ KUNZLER, Tabelião de Notas, sobre cuja identidade e capacidade jurídica dou fé. E, pelo outorgante foi dito que, nomeia e constitui seu bastante procurador, RODRIGO ALVES SOARES, brasileiro, gerente comercial, portador da carteira nacional de habilitação sob número 01030737444, expedida pelo DETRAN/RS em 19/01/2022, inscrito no CPF sob número 481.149.520-91, filho de Cícero Dias Soares e Carolina Alves Soares, casado, residente e domiciliado na Rua Guerino Zugno, nº 1476, Bairro Samuara, na cidade de Caxias do Sul/RS., a quem confere poderes para representar a empresa outorgante, bem como as suas filiais: Filial nº 01, devidamente registrada na Junta Comercial do Estado do Rio Grande do Sul, sob NIRE 42.901.114.281, inscrita no CNPJ 02.543.216/0006-33, com estabelecimento na Avenida Rio Branco, nº 404, Torre II, Sala 908-A, Centro, no Município de Florianópolis/SC., Filial nº 02, devidamente registrada na Junta Comercial do Estado do Rio Grande do Sul, sob NIRE 53.900.356.328, inscrita no CNPJ 02.543.216.0007-14, com estabelecimento na SCS Quadra 06, Bloco A, nº 130, 6º andar, parte H2, Edifício Ermes, Bairro Asa Sul, no Município de Brasilia/DF., Filial nº 03, devidamente registrada na Junta Comercial do Estado do Rio Grande do Sul, sob NIRE 41.901.686.313, inscrita no CNPJ 02.543.216/0008-03, com estabelecimento na Rua Antônio Lacerda Braga, nº 960, Sala J, Cidade Industrial Curitiba, no Município de Curitiba/PR., Filial nº 04, devidamente registrada na Junta Comercial do Estado do Rio Grande do Sul, sob NIRE 35.905.941.194, inscrita no CNPJ 02.543.216/0010-10, com estabelecimento na Rua Alameda Rio Negro, nº 1030, Escritório 206, Sala 03, Condomínio Stadium, Bairro Alphaville Centro Industrial e Empresarial, no Município de Barueri/SP., Filial nº 05, devidamente registrada na Junta Comercial do Estado do Rio Grande do Sul, sob NIRE 32900792252, inscrita no CNPJ 02.543.216/0011-09, com estabelecimento na Avenida Acesso Rodoviário, s/nº, quadra 06, sala 107, Bairro Terminal Intermodal da Serra, no Município de Serra/ES., Filial nº 06, devidamente registrada na Junta Comercial do Estado do Rio Grande do Sul, sob NIRE 52901653261, inscrita no CNPJ 02.543.216/0012-81, com estabelecimento na Avenida 136, nº 761, sala B142, Bairro Setor Sul, no Município de Goiânia/GO., podendo para tanto: A - representar em processos licitatórios em toda a esfera nacional: órgãos públicos municipais, estaduais e federais, Conselhos Regionais em todo o território Nacional, bem como representa-la junto aos órgãos componentes do Sistema S, podendo assinar propostas, declarações, atestados, contratos, recursos administrativos, credenciamentos, enfim, todas as práticas comerciais envolvidas no processo; B - representar em atividades comerciais com empresas de capital misto e privado, cooperativas, em toda a esfera nacional; C assinatura como representante legal, podendo substabelecer poderes em processos licitatórios; D representar perante Sindicatos e Ministério do Trabalho, quando em questões referentes à contratação e demissão de funcionários, podendo enfim, assinar e requerer tudo o que for necessário



para o cumprimento do presente mandato. (LAVRADA SOB MINUTA). Fica reservado à outorgante, o uso simultâneo dos poderes ora conferidos. O nome e dados do outorgado, bem como os elementos relativos ao objeto do presente instrumento, são de inteira responsabilidade da outorgante e foram por este conferidos, uma vez que não apresentou os documentos pertinentes, devendo fazê-lo por ocasião da utilização deste instrumento. E, assim me pediu que lhe lavrasse a presente procuração que, após lida e achada conforme, aceita, ratifica e assina. Foram dispensadas as testemunhas instrumentárias. Ana Paula Klering, Escrevente autorizada, a digitou e eu DAICIR JOSÉ KUNZLER, Tabelião de Notas, assino em público e raso. Consulte a autenticidade deste ato acessando https://balcaodigital.skyinformatica.com.br/ informando a chave de acesso 24Y53DFCB e o digito validador 040. Emol. Escr. s/ cont. financeiro: R\$ 88,80 (0215.04.1900002.10150 = R\$ 4,40); Processamento eletrônico: R\$ 6,00 (0215.01.2200005.76665 = R\$ 1,80)

FARROUPILHA, SEXTA-FEIRA, 30 DE JUNHO DE 2023.

Assinado digitalmente por: DAICIR JOSE KUNZLER CPF: 007.742.510-34 Certificado emitido por AC SAFEWEB RFB v5 Data: 30/06/2023 11:38:46 -03:00



DAICIR JOSÉ KUNZLER Tabelião de Notas Assinado Digitalmente



A consulta estará disponível em até 24h no site do Tribunal de Justiça do RS http://go.tjrs.jus.br/selodigital/consulta Chave de autenticidade para consulta 103796 51 2023 00043282 91





Sistema de Cadastramento Unificado de Fornecedores - SICAF

Declaração

Declaramos para os fins exigidos na legislação, conforme documentação registrada no SICAF, que a situação do fornecedor no momento é a seguinte:

Dados do Fornecedor

CNPJ:	02.543.216/0011-09	DUNS®: 92****71
Razão Social:	PERFIL COMPUTACIO	NAL LTDA
Nome Fantasia:		
Situação do Fornecedor:	Credenciado	Data de Vencimento do Cadastro: 28/04/2025
Natureza Jurídica:	SOCIEDADE EMPRESÁ	ARIA LIMITADA
MEI:	Não	
Porte da Empresa:	Demais	

Ocorrências e Impedimentos

Ocorrência:	Consta
Impedimento de Licitar:	Nada Consta

Níveis cadastrados:

Automática: a certidão foi obtida através de integração direta com o sistema emissor. Manual: a certidão foi inserida manualmente pelo fornecedor.

I - Credenciamento

II - Habilitação Juridica

III - Regularidade Fiscal e Trabalhista Federal

Receita Federal e PGFN FGTS Trabalhista (http://www.tst.jus.br/certidao)	Validade: Validade: Validade:	14/09/2025 01/04/2025 14/09/2025	Automática Automática Automática
IV - Regularidade Fiscal Estadual/Distrital e	Municipal	, ,	
Receita Estadual/Distrital	Validade:	26/05/2025	
Receita Municipal	Validade:	10/05/2025	
VI - Qualificação Econômico-Financeira			
	Validade:	31/05/2025	

Esta declaração é uma simples consulta e não tem efeito legal

Emitido em: 18/03/2025 14:04 CPF: 805.XXX.XXX-49 Nome: IGOR SIDNEI REOLON Ass: _____ 1 de 1



Cyber Resilient Security in Dell PowerEdge Servers

October 2023 H19738.1

White Paper

Abstract

This white paper highlights the Dell PowerEdge Cyber Resilient Architecture and describes the server life cycle for implementing zero trust principles for your infrastructure. Dell PowerEdge security controls provide a comprehensive security solution that ensures resiliency while enforcing a zero-trust posture.

D

L

Technologies



Copyright

The information in this publication is provided as is. Dell Inc. makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Use, copying, and distribution of any software described in this publication requires an applicable software license.

Copyright © 2023 Dell Inc. or its subsidiaries. Published in the USA October 2023 H19738.1.

Dell Inc. believes the information in this document is accurate as of its publication date. The information is subject to change without notice.



Contents

Contents

Executive summary	5
Overview	5
Revisions	5
We value your feedback	5
Introduction	6
Digital infractructure complexity	6
Sophistication and complexity of threats	
Begulatory landscape and internal mandates	
Zero Trust strategy for the modern world	
Core principles of Zero Trust	0
Seven pillars of Zero Trust	7 Q
The Dell advantage	0
Security journey across the server life cycle	
First stage – Choosing the server	10
Challenges	10
PowerEdge solutions	10
Dell Secure Development Lifecycle	11
Compliance advantage	11
Rapid response to new vulnerabilities	12
Bug Bounty Program	12
Solutions covering threat vectors for every layer of the server	13
Occurred atoms - Overally shain a consister	
Second stage – Supply chain security	
FowerEage solutions	
End-to-end supply chain assurance	
Secured Component Verification	
Software Bill of Materials	
Third stage – Efficient deployment and configuration at-scale	17
Challenges	17
PowerEdge solution	17
System integrity	17
Hardware security	20
Protecting data at rest	21
Protecting Data in Flight	23
Protecting data in use	26
Identity Access Management	28
Capabilities and automation for efficient at-scale deployment	30



Contents

Fourth Stage – Security management and monitoring	
Challenges	
PowerEdge Solutions	
Visibility, logging, and alerts	
SELinux framework	
Real-time detection – BIOS Live scan	
Silicon-based Root of Trust	
Automated and manual recovery	
Updating	
Restoring server configuration after hardware servicing	
CloudIQ	
Managed detection and response services	
Fifth stage – Secure decommissioning and repurposing	40
Challenge	40
PowerEdge solutions	40
Secure Erase	41
Secure erase – physical disk	41
Data sanitization and destruction services	42
Summary	42
References	44
Dell Technologies documentation	



Executive summary

Overview

Revisions

The Dell Technologies approach to security is intrinsic in nature – it is built-in, not boltedon later, and it is integrated into every step of Dell's Secure Development Lifecycle. We continuously strive to evolve our Dell PowerEdge security controls, features, and solutions to meet the ever-changing threat landscape and to help customers accelerate Zero Trust adoption.

Securing your infrastructure is not a one-time investment, but a mindset and overall approach. This white paper uses this journey perspective to describe the PowerEdge advantages across the server life cycle. In each of the life cycle phases, from deployment through maintenance to decommissioning, we highlight how Dell's PowerEdge Cyber Resilient Architecture security features work together to provide both resiliency and a zero-trust approach. The Dell Remote Access Controller (iDRAC9) enables many of these features.

We continue to anchor security with a Silicon-based Root-of-Trust (RoT). Since the previous PowerEdge cyber resilient security whitepaper, many new features have been added that span from access control to data encryption to supply chain assurance. All features make extensive use of intelligence and automation to help you stay ahead of the threat curve, and to enable the scaling demanded by ever-expanding use models. Dell's Cyber Resilient Architecture, enhanced over many years, is the foundation for the critical elements of a Zero Trust environment.

Date	Part number/ revision	Description
November 2022	H19738	Initial release
October 2023	H19738.1	Updated to include the 16 th generation of PowerEdge servers

We value yourDell Technologies and the authors of this document welcome your feedback on this
document. Contact the Dell Technologies team by email.

Authors: Deepak Rangaraj, Kim Kinahan

Contributors: Marshal Savage

Note: For links to other documentation for this topic, the <u>Dell Technologies Info Hub for</u> <u>PowerEdge</u>.



Introduction

Digital infrastructure complexity	Modern IT environments have changed drastically in the past few years with servers being deployed in various use cases such as on-premises, multicloud, Edge, Telco, and so on. The server platforms are becoming more complex with an ever-increasing number of components that require firmware for configuration and management. We are generating data at a speed and volume higher than ever and this data is often generated and stored at numerous locations distributed geographically. This increasing complexity necessitates effective management of security controls to mitigate the expansion of the attack surface.
Sophistication and complexity of threats	The Dell Technologies Digital Transformation Index found that data privacy and cybersecurity concerns are the leading barriers to digital transformation ¹ . The complexity, sophistication, and frequency of cyberattacks are increasing, and the damage caused by attacks is becoming more costly. Complicating matters further, today's threat actors are taking advantage of technological advancements, such as AI and a lower cost of entry. Malicious actors are continuously searching for vulnerabilities to exploit. With the assistance of advanced AI systems, they can carry out nefarious activities at an unprecedented scale and manipulate systems in innovative and harmful ways beyond human capabilities. Global damages related to cybercrime are predicted to reach \$10.5 trillion by 2025. ²
Regulatory landscape and internal mandates	As global threats increase, governments worldwide are developing regulatory guidance in response to cyber threats. As a result, private institutions are creating stronger policies and mandates to mitigate advanced persistent threats. Also, there are more mandates for security requirements that are needed to work with the government. These requirements impact suppliers, vendors, and any organization that partners with the government. In addition, regulations are carrying over to the critical infrastructure sectors, such as healthcare, transport, and finance. Outside of government regulations, many customers want to harden their infrastructure and are developing their own internal mandates or security policies.
Zero Trust strategy for the modern world	The increasing infrastructure complexity and threat landscape are driving a critical need to secure not only infrastructure hardware, but also the firmware and the supply chain itself. When applying a Zero Trust strategy, customers focus on business controls, the control plane, and applications and data. However, there is a critical need to secure what is below these items, including infrastructure hardware and firmware, supply chain for the infrastructure, and design and processes used to build the infrastructure. Zero Trust principles must be applied to all these aspects for more comprehensive cyber resilience.
	Dell Technologies has security built into our industry-leading servers, storage, HCI, and data protection appliances to help protect data wherever it is stored, managed, or used. As a foundation for securing our PowerEdge server products, they are foremost cyber resilient – capable of anticipating, withstanding, and recovering from cyber threats.
	¹ Dell Technologies 2020 Digital Transformation Index
	² www.cybersecurityventures.com /cybercrime-damage-costs-10-trillion-by-2025/

293

Simultaneously, the process of designing PowerEdge security controls and tools incorporates Zero Trust principles. By anticipating how customers want to use these capabilities while they are setting up their Zero Trust deployments, we have adapted our approach. We made it easier to work with us no matter where customers are on their journey towards Zero Trust adoption.



Figure 1. Infrastructure security and Zero Trust

The Dell infrastructure security approach is integral to Dell products and accelerates Zero Trust adoption because it is:

- Designed and built using Zero Trust principles
- Provides Zero Trust capabilities and features
- Provides common, consistent behavior

Core principles of Zero Trust

The tenets of Zero Trust architecture are built on a set of principles that presume that the *network is always vulnerable to compromise* and sets out to safeguard access to critical data and resources. Unlike trust-then-verify frameworks, a zero-trust approach eliminates implicit trust. Every user, device, and application must be continuously *authenticated and explicitly authorized* based on a range of factors such as identity, device status, location, and behavior.

Identity plays a crucial role in Zero Trust. Identification pertains not only to people but to applications, communication paths, network devices, and the data itself. When an IT asset is identified, authenticated, and explicitly authorized, the principle of *least privilege* is applied. This approach ensures that only authorized entities are given a minimum level of access required to perform their specific task. The data-centric security model constantly limits access while also looking for anomalous or malicious activity. The Zero Trust approach reduces the granularity of validation at key intersections for verified trust, optimizing least privilege without impacting workload efficiency. The goal is to deter attacks and reject them at point of entry. However, if a breach occurs, the amount of damage is minimized, along with an enhanced ability to detect and remediate immediately.

Adopting the Zero Trust mindset and using Zero Trust principles enables systems administrators to control how users, processes, and devices engage with data. These



Introduction

principles can prevent the abuse of compromised user credentials, remote exploitation, insider threats, and even mitigate the effects of malicious supply chain activity.

Seven pillars of Zero Trust The Zero Trust model, defined by the National Institute of Standards and Technology (NIST)³, identifies seven interrelated pillars that work together to provide a comprehensive and holistic approach to infrastructure and data security. Each pillar represents a specific functional or key focus area for implementation of Zero Trust security controls. When combined, these seven pillars provide a multifaceted, layered, and integrated security framework.

Dell's Zero Trust approach integrates a broad set of security controls and automation capabilities for the management of the infrastructure and the applications running on it. The following table highlights Dell's capabilities across the seven pillars as outlined by NIST:

Zero Trust pillar	NIST description ³	PowerEdge highlights
User	 User identification, authentication, and access control: Only validated and authorized users can access data and resources. The principle of <i>least privilege</i> is applied where users are granted the minimum level of access required to perform their specific tasks. 	 Identity and Access Management Multi-Factor Authentication—RSA Secure ID Active Directory or LDAP integration with single sign-on (SSO) support Role-based access control and auditing
Device	 Monitoring and enforcement of device health, compliance, and device posture assessment: Monitoring - looking for anomalies and suspicious read/write activity. Health – confirming the latest version of the firmware. All devices are identified, inventoried, authorized, authenticated, and updated. 	 Silicon Root-of-Trust (ROT) with complementary Intel Boot Guard and AMD Platform Secure Boot (AMD PSB) Secure supply chain with Secured Component Verification (SCV) Chassis locks and intrusion detection Dynamic USB port enable/disable Trusted Platform Module (TPM) Device attestation with SPDM (Security Protocol Data Model from DMTF)
Data	Ensure data transparency and visibility by using enterprise infrastructure, applications, standards, solid end-to-end encryption, and data tagging.	 Data-at-rest protection: Drive encryption with local (LKM) or Secure Enterprise Key Management (SEKM) with direct-attached NVMe drive support Baseboard Management Controller (BMC)-based Local Key Management (iLKM) Data-in-use protection:

Table 1. Dell Zero Trust approach across the seven pillars

³ National Institute of Standards and Technology (NIST) Special Publication (SP) 800-207, Zero Trust Architecture

8 Cyber Resilient Security in Dell PowerEdge Servers



Introduction

Zero Trust pillar	NIST description ³	PowerEdge highlights
		Confidential compute—Intel SGX, Intel MKTME, AMD Secure Memory Encryption (SME) AMD SME, AMD Secure Encrypted Virtualization (SEV) SEV, AMD persistent memory encryption
Application and Workload	Secure applications and workloads, and protect containers and VMs.	 Secure Development Lifecycle Cryptographically signed BIOS and firmware updates Secure end-to-end boot and Unified Extensible Firmware Interface (UEFI) boot capabilities Drift detection Rapid Response and mitigations for CVEs
Network and Environment	Encrypt, monitor, and analyze network. Logically and physically segment, isolate, and control the network and the environment (on- premises and off-premises) using granular access and policy restrictions.	 Dedicated BMC (iDRAC) network module SSH/TLS communication options TLS 1.3 support DPU/SmartNIC
Visibility and Analytics	Monitor activities and behaviors across the infrastructure (user, device, data, network, and application) to identify patterns and anomalies. Use analytics to detect and respond to security threats.	 Persistent event logging and auditing Real-time and boot time firmware scanning Security alerts CloudIQ
Automation and Orchestration	Automate manual security and other applicable processes to take policy-based actions across the enterprise with speed and at scale.	 OpenManage Enterprise drift detection Firmware rollback Automatic BIOS and operating system recovery Centralized updates Automatic SSL certificate renewal

The Dell advantage

Security is in our DNA, and we are committed to making our products secure by design and secure by default. PowerEdge servers are built using Zero Trust principles internally and have capabilities that enable customers to set up a Zero Trust IT environment and operations. Our products also strive to provide common and consistent behavior and controls across our portfolio.

PowerEdge servers with iDRAC9 have an integrated immutable silicon-based platform Root-of-Trust (RoT) that is used to establish a verified chain of trust that extends throughout the server life cycle, from deployment through maintenance to decommissioning. This RoT combined with security controls and comprehensive management tools provides robust layers of security across the PowerEdge hardware and firmware.

Cyber Resilient Security in Dell PowerEdge Servers 9



First stage – Choosing the server

These PowerEdge capabilities not only ensure cyber resiliency to protect, detect, and recover from attacks, they also maintain a locked-down posture for a Zero Trust approach of least privilege. Least privilege ensures that users and devices are only given access to what they need to perform their tasks. Our goal is to make Zero Trust a reality for our customers and accelerate its speed of adoption.

Security journey across the server life cycle

Implementing security across your infrastructure requires a series of ongoing efforts and measures to protect underlying systems, networks, and resources. Transitioning to a mature security model is not a one-time investment and cannot be accomplished overnight. It is an ongoing journey and approach of implementing stringent security policies. As your Zero Trust implementation matures over time, enhanced visibility and extensive controls allow you to keep pace with the threat landscape. We have organized this journey into five stages.

- **Stage 1: Choosing the server**—Customers want assurance that security is a top priority and is built into every aspect of design.
- Stage 2: Ensuring supply chain security—Customers face the real risk of malicious offenders replacing original components with counterfeits, implants, or malware.
- Stage 3: Efficient deployment and configuration at scale—How a server is deployed has a direct impact on its performance, stability, and security. Proper planning and configuration are essential to ensure that the server is set up correctly and that all necessary components are in place.
- Stage 4: Security management and monitoring— Because attacks happen quickly, faster than a human can detect, customers must proactively monitor their environment to take quick action. The lack of skills and training can also exacerbate the problem.
- Stage 5: Secure repurposing and decommissioning—Data security is a key consideration when the server is repurposed or retired. IT best-practices recommend removing all data from the server to ensure that confidential information is not inadvertently shared or compromised.

First stage – Choosing the server

Challenges Customers want to be confident that their server is secure in its entirety and does not introduce any vulnerabilities into their environment. They want assurances from server vendors addressing every aspect of design, the entire supply chain, including hardware, software, and firmware. This level of assurance differentiates on-premises infrastructure when compared to cloud service providers who might only offer a black box infrastructure, leaving customers uncertain about the security of the underlying components of their cloud-based system.

PowerEdge
solutionsThe first step of the Zero Trust journey with Dell Technologies starts before you receive
your PowerEdge server. Intrinsic security practices are incorporated into hardware
product design and software or firmware code development. These practices include
processes and policies that ensure security features are implemented at the time of
product inception and continue throughout the development cycle. In essence, security is

10 Cyber Resilient Security in Dell PowerEdge Servers

ls. 297



Dell Secure Development Lifecycle

Delivering the Cyber Resilient Architecture requires security awareness and discipline at each stage of development. The Secure Development Lifecycle (SDL) model is used to make security an integral part of the overall server design process. Some key aspects of this process include:

- Features that are conceived, designed, prototyped, implemented, set into production, deployed, and maintained with security as a key priority
- Server firmware that is designed to obstruct, oppose, and counter the injection of malicious code during all phases of the product development life cycle:
 - Threat modeling and penetration testing provide coverage during the design process.
 - Secure coding practices are applied at each stage of firmware development.
- For critical technologies, external audits that supplement the internal SDL process to ensure that firmware adheres to known security best practices
- On-going testing and evaluation of new potential vulnerabilities using the latest security assessment tools
- Rapid response to critical Common Vulnerabilities and Exposures (CVEs) including recommended remediation measures if warranted



Figure 2. Dell Secure Development Lifecycle

Compliance advantage

Dell Technologies has received the certifications that are needed to comply with major United States Federal and other global governmental requirements as well as industry standards such as from NIST, as described in the following table:



First stage - Choosing the server

Table 2.Certification descriptions

Certification	Description
Common Criteria	Certain configurations of PowerEdge Servers include components with common criteria certifications. (For example, IDRAC and TPM)
FIPS 140	Certain configurations of PowerEdge Servers include FIPS 140-certified cryptographic modules (For example, TPM, IDRAC9, Chassis Management Controller (CMC), Self-Encrypting Drives (SEDs), and SSDs)
IPv6	 PowerEdge Servers are fully USGv6r1 and IPv6 Ready Logo- compliant and certified with IPv6-only capabilities while running operating systems such as Red Hat Enterprise Linux 8.4 and the applicable versions of Windows 2019 or Windows 2022 Server.
	 Dell PowerEdge iDRAC9 (with 5.1x firmware) is USGv6r1 and IPv6 Ready Logo-compliant and certified with IPv6-only capabilities.

Rapid response to new vulnerabilities

CVEs are newly discovered attack vectors that compromise software and hardware products. Timely responses to CVEs are critical to most companies so they can swiftly assess their exposure and take appropriate action.

Dell Technologies works aggressively to respond quickly to new CVEs in our PowerEdge servers and provide timely information including the following:

- The products that are affected
- Remediation steps
- If needed, availability of updates to address the CVEs

Bug Bounty Program

Dell Technologies recognizes the value of the security research community to broaden visibility into potential vulnerabilities and threats and welcomes the opportunity to collaborate with community members who share this common goal.

Dell's Bug Bounty Program applies to security vulnerabilities identified in Dell-branded or currently supported products.

Solutions covering threat vectors for every layer of the server

There are many threat vectors in today's changing landscape. The following tables summarize the Dell approach to managing critical threats in each of the server layers.

Server platform layers		
Security layer	Threat vector	Dell solution
Physical server	Server/component tampering or theft of component	 Secured Component Verification (SCV) Chassis Intrusion Detection Secure Enterprise Key Management Intel TME SPDM
Firmware and software	Firmware corruptionMalware injection	 Silicon-based Root-of-Trust Intel Boot Guard AMD Secure Root-of-Trust UEFI Secure Boot Customization Cryptographically signed and validated firmware
	Software	CVE reportingPatching as required
Attestation trust features	Server identity spoofing	 TPM Intel TXT Chain of trust 802.1x features SPDM
Server management	Rogue configuration and updatesUnauthorized open-port attacks	iDRAC9 Remote attestation

 Table 3.
 Dell solution for common server platform layer threat vectors



Second stage - Supply chain security

Server environment layers		
Security layer	Threat vector	Dell Technologies solution
Data	Data breach	 Self-Encrypting Drives (SED) – FIPS or Opal/TCG Secure Enterprise Key Management ISE-only (Instant Secure Erase) drives Secure User Authentication
Supply Chain Integrity	Counterfeit componentsMalware threats	 ISO9001 certification for all global server manufacturing sites Secured Component Verification Proof of possession Software Bill of Materials (SBOM) Security measures implemented as part of the Secure Development Lifecycle
Supply Chain Security	 Physical security in manufacturing sites Theft and tempering during transport 	 Transported Asset Protection Association (TAPA) facility security requirements Customs-Trade Partnership Against Terrorism (C-TPAT) Secured Component Verification

Table 4. Dell solution for common server environment layer threat vectors

Second stage – Supply chain security

Challenges

Modern server platforms are becoming more complex with hundreds of components that require firmware for configuration and management. As a result, the server supply chain is also becoming increasingly complex with hundreds of third-party vendors supplying components as well as the use of open-source software. This complex supply chain contributes to an increase in the attack surface available for threat actors if not managed properly. Customers face the risk of vulnerabilities and threats being introduced into their environments if the supply chain integrity is not assured. The two main aspects of supply chain integrity include:

- **Maintaining hardware integrity**—Ensuring that there is no product tampering or the insertion of counterfeit components or malicious implants before the product is shipped to customers.
- **Maintaining software integrity**—Ensuring that no malware is inserted in firmware or device drivers before shipping the product to customers and preventing code with known vulnerabilities from being introduced into the environment.

inventory, information, intellectual property, and people. These security measures provide

 PowerEdge solutions
 End-to-end supply chain assurance

 Dell Technologies employs a multifaceted approach to protect its supply chain and delivers solutions that customers can trust in an increased threat environment. Our supply chain security consists of prevention and detection controls that protect physical assets,

14 Cyber Resilient Security in Dell PowerEdge Servers

Second stage – Supply chain security

supply chain assurance and ensure integrity by reducing opportunities for the malicious or negligent introduction of malware and counterfeit components into the supply chain.

Dell supply chain controls span supplier selection, sourcing, production processes, and governance through auditing and testing. When a supplier has been selected, the new product introduction process verifies that all materials used during all build stages are sourced from the approved vendor list and match the bill of materials as appropriate. Material inspections during production help identify components that are mismarked, deviate from normal performance parameters, or contain an incorrect electronic identifier.

Parts are procured directly from the Original Design Manufacturer (ODM) or Original Component Manufacturer (OCM) when possible. The material inspection that occurs during the new product introduction process provides multiple opportunities to identify counterfeit or corrupted components that might have entered the supply chain.

Additionally, Dell Technologies maintains ISO 9001 certification for all global manufacturing sites. Strict adherence to these processes and controls helps minimize the risk of counterfeit components being embedded among the Dell products or malware being inserted into firmware or device drivers. As part of SDL, Dell Technologies has several long-standing, key practices that establish and maintain security in manufacturing facilities and logistical networks.

Facilities used to design, build, customize, or fulfill Dell products must demonstrate compliance with several internationally recognized physical security standards such as those defined by the Transported Asset Protection Association (TAPA), American Society for Industrial Security (ASIS), International Standards Organization (ISO), and the Business Alliance for Secure Commerce (BASC).

Protective measures have also been put in place to guard products against theft and tampering during transport as part of an industry-leading logistics program. This program provides a continuously staffed command center to monitor select inbound and outbound shipments across the globe to ensure that shipments make it from one destination to another without disruption.

Dell Technologies audits suppliers and facilities, addressing various factors, including the use of digital closed-circuit TV cameras, access control systems, intrusion detection, and guard service protocols. Other controls are applied to protect Dell cargo during the shipping and logistics process, including tamper-evident packaging, cargo locks and seals, and threat intelligence monitoring of key freight lanes. Internet of Things (IoT) tracking devices are also deployed on select shipments to enable real-time telemetry data monitoring to escalate any security noncompliance events observed during transit.

Dell Technologies also maintains certifications in multiple secure trade and commerce programs such as Tier 3 status with the United States Customs and Border Protection's Customs Trade Partnership Against Terrorism (C-TPAT), Canada's Partners in Protection (PIP), Singapore's Secure Trade Partnership, and Authorized Economic Operator (AEO) status in several other nations. These programs are internationally recognized by member states of the World Customs Organization and demonstrate "best in class" supply chain security standards in the private sector. These programs focus on supplier accountability, security management policies, counter smuggling, trafficking controls, and tamper prevention – all intended to secure trade across international borders.



Second stage - Supply chain security

Supply chain integrity ensures that customers' products are safely delivered and when received, operate as intended. An important feature of supply chain integrity is the development of hardware and software baseline specifications that are preserved securely and later used as a reference to verify that no unauthorized modifications have been made.

Secured Component Verification

Dell Technologies' Secured Component Verification (SCV) for PowerEdge is a supply chain assurance offering that verifies that the server received by a customer matches the configuration that was shipped from the factory. The factory generates a certificate that contains unique component IDs for a specific server. This certificate is stored in a cryptographically secure vault in iDRAC. On receiving the server, the customer runs the SCV application on the host to generate an inventory of the current system, including unique component IDs, and then validates it against the golden factory inventory in the SCV certificate stored in iDRAC.

The SCV application generates a report that identifies any component mismatches from what was installed in the factory. It also verifies the certificate and Chain of Trust along with the Proof of Possession of the SCV Private key for iDRAC. The current implementation supports direct ship customers and does not include Value Added Reseller (VAR) or Part Replacement scenarios.





Software Bill of Materials

As part of Dell's software supply chain security controls, NIST standards, and in alignment with the President's Executive Order (EO) 14028, a Software Bill of Materials (SBOM) is available for a limited number of products across our portfolio. Dell SBOM data adheres to the Software Package Data Exchange (SPDX) standard and is provided in JSON format. SBOM data provides software supply chain transparency and can be used in vulnerability scanning and asset tracking tools on the customer side.

SBOM enables customers to gain a clearer understanding of the software components, versions, licenses, and any open-source software used on the platform. It can facilitate faster detection of known security vulnerabilities in the software components, ultimately enhancing security.

Third stage – Efficient deployment and configuration at-scale

Challenges

Cyber resilient configuration and deployment of the servers is a critical step in the server life cycle. Any mistakes or oversights during this process can result in poor server performance, downtime, or even security breaches. The right set of controls and gates must be in place to ensure system and data integrity throughout the server operation. This configuration and deployment must be performed for hundreds or thousands of servers while ensuring consistency and minimizing any manual errors.

PowerEdge

solution

System integrity

Ensuring system integrity is foundational for securing the server and establishing a lockdown posture for Zero Trust operations. This integrity starts with ensuring that server hardware and components are genuine, and from a trusted and authorized source. Then, the firmware and software must be verified to ensure that a bad actor has not tampered with them. For PowerEdge, this process of ensuring system integrity starts with a siliconbased platform RoT. This RoT anchors the other security controls on the server platform and establishes a chain of trust for cryptographic verification of hardware and software components on the server.

RoT as the anchor for cryptographically verified trusted booting

One of the most critical aspects of server security is ensuring that the boot process can be verified as secure. This process provides a trusted anchor for all subsequent operations such as booting an operating system or updating firmware.

All PowerEdge servers have an immutable, silicon-based RoT burnt into them from the factory. The RoT has one-time programmable, read-only public keys that can be used for cryptographic verification and attestation of integrity.

The BIOS boot process uses Intel Boot Guard technology or AMD PSB technology that cryptographically verifies the BIOS code to be loaded. A verification failure results in a server shutdown and the Lifecycle Controller Log includes a notification. The IT administrator can then initiate the BIOS recovery. If Boot Guard validates successfully, a chain of trust procedure validates the remaining BIOS modules until control is handed off to the operating system or hypervisor.

In addition to Boot Guard's verification mechanism, iDRAC9 4.10.10.10 or higher provides a RoT mechanism to verify the BIOS image at the host boot time. The host can boot only after the BIOS image is successfully validated. iDRAC9 also provides a mechanism to validate the BIOS image at run time, on demand, or at user-scheduled intervals.

Cryptographically verified Trusted Booting

One of the most critical aspects of server security is ensuring that the boot process can be verified as secure. This process provides a trusted anchor for all subsequent operations such as booting an operating system or updating firmware.

PowerEdge servers have used silicon-based security for several generations for features such as iDRAC Credential Vault, an encrypted secure memory in iDRAC for storing sensitive data. The boot process is verified using a silicon-based RoT to meet recommendations in NIST SP 800-147B (BIOS Protection Guidelines for Servers) and NIST SP 800-155 (BIOS Integrity Measurement Guidelines).



Third stage - Efficient deployment and configuration at-scale

Security Protocol Data Model for component attestation

The Distributed Management Task Force (DMTF), of which Dell Technologies is a leading member, defines the Security Protocol Data Model (SPDM). SPDM defines a consistent, open-standard method of communicating in the server to gather information about server components. This component information is protected by encrypting it and using authenticated key exchange to integrity-protect all communications between components. iDRAC SPDM implementation provides visibility into the PERC12 and certain NIC components. As part of the hardware inventory, iDRAC verifies the authenticity and integrity of PERC12 and NIC devices by cryptographically verifying the identity, firmware, and configuration.

TPM Support

PowerEdge servers support two versions of TPM:

- TPM 2.0 FIPS + Common Criteria + TCG certified (Nuvoton)
- TPM 2.0 China (NationZ)

TPM can be used to perform public key cryptographic functions, compute cryptographic hash functions, generate, manage, and securely store keys, and perform attestation. Intel's Trusted Execution Technology (TXT) functionality and Microsoft's Platform Assurance feature in Windows Server 2016 is also supported. TPM can also be used to enable the BitLocker hard drive encryption feature in Windows Server 2012, Windows Server 2016, and Windows Server 2022.

Attestation and remote attestation solutions can use TPM to take measurements at boot time of a server's hardware, hypervisor, BIOS, and operating system, and compare them in a cryptographically secure manner against "golden" or "base" measurements. These measurements are commonly stored outside the TPM on a remote attestation server solution. The TPM PCR measurements stored in the TPM are recalculated on every boot. If they are not identical, the server system might have been compromised and system administrators can disable and disconnect the server either locally or remotely.

Servers can be ordered with or without TPM, but for many operating systems and other security provisions, it is becoming a standard. TPM is enabled through a BIOS option. It is a Plug-In Module solution; the planar has a connector for this plug-in module.

UEFI Secure boot for firmware

PowerEdge servers also support industry-standard UEFI Secure Boot that checks the cryptographic signatures of UEFI drivers and other code loaded before the operating system is running. Secure Boot represents an industry-wide standard for security in the preboot environment. Computer system vendors, expansion card vendors, and operating system providers collaborate on this specification to promote interoperability.

When enabled, UEFI Secure Boot prevents unsigned (that is, untrusted) UEFI device drivers from being loaded, displays an error message, and does not allow the device to function. You must disable Secure Boot to load unsigned device drivers.

In addition, the 14th, 15th, and 16th generations of PowerEdge servers offer the unique flexibility of using a customized boot loader certificate. This certificate is primarily a feature for administrators of Linux environments that want to sign their own operating system boot loaders instead of relying on the default signing certificate provided by Microsoft's UEFI Certificate Authority (CA). Custom certificates can be uploaded by using the preferred

18 Cyber Resilient Security in Dell PowerEdge Servers

Third stage - Efficient deployment and configuration at-scale

iDRAC API to authenticate the customer's specific operating system boot loader. The NSA cites this PowerEdge UEFI customization method for mitigating Grub 2 vulnerabilities in servers.⁴ PowerEdge supports complete customization of Secure Boot, including removal of all industry-standard certificates provided by Microsoft, VMware, or the UEFI CA.

Intel Boot Guard and AMD PSB

Intel Boot Guard and AMD PSB are host processor features that provide strong firmware integrity guarantees, by preventing firmware that is not authorized by the Dell OEM from running on the system. By enabling these features as additional defense-in-depth measures, certain classes of physical attacks are mitigated, such as flash memory replacement or reprogramming, and Time-of-Check-Time-of-Use (TOCTOU) race conditions. All combined, the RoT features in the system make compromise of the trusted computing base (TCB) difficult.

iDRAC/BMC

The Integrated Dell Remote Access Controller (iDRAC) is a Baseboard Management Controller (BMC) that is integrated in Dell PowerEdge servers. iDRAC provides secure and remote server access for many common management functions; administrators can deploy, manage, monitor, update, troubleshoot, and remediate Dell servers from any location without the use of agents and out of band.

iDRAC offers industry-leading security features that adhere to and are certified against well-known NIST standards, Common Criteria, and FIPS 140-2. It is through iDRAC9 that the end user can configure security features to maximize the security posture of the system.

iDRAC credential vault

The iDRAC service processor provides a secure storage memory that protects sensitive data such as iDRAC user credentials and private keys for self-signed SSL certificates. Another example of silicon-based security, this memory is encrypted with a unique immutable root key that is programmed into each iDRAC chip at the time of manufacture. This memory protects against physical attacks where the attacker desolders the chip to gain access to the data.

SELinux framework

SELinux operates at the core kernel level on iDRAC and does not require user input or configuration. SELinux logs security messages when an attack is detected. These log messages indicate when and how an attacker tried to break into the system. These logs are available through SupportAssist to enrolled customers. In future iDRAC releases, these logs will be available in the Lifecycle Controller Logs.

Factory-generated default passwords

By default, all 14th, 15th, and 16th generations of PowerEdge servers ship with a unique, factory-generated iDRAC password to provide additional security. This password is on the pull-out Service Tag on the front of the chassis, next to the server asset label. If you choose this default option, use this password to log in to iDRAC for the first time, rather than using a universal default password. For security purposes, Dell Technologies strongly recommends changing the default password.

⁴ <u>CTR-UEFI-Secure-Boot-Customization-UOO168873-20.PDF (defense.gov)</u>



Third stage - Efficient deployment and configuration at-scale

Hardware security

Hardware security is an integral part of any comprehensive security solution. Some customers want to limit access to external ports, such as USB. Generally, a server chassis does not need to be opened after it has been put into production. At a minimum, customers always want to track and log any such activities. The overall goal is to discourage and limit any physical intrusion.

Chassis intrusion detection and alert

PowerEdge servers provide hardware intrusion detection and logging, with detection working even when no AC power is available. Sensors on the chassis detect when anyone opens or tampers with the chassis, even during transit. Servers that have been opened while in transit generate an entry in the iDRAC Lifecycle log after power is supplied.

Dynamic USB port management

For more security, you can completely disable USB ports. You also have the option of disabling only the USB ports on the front. For example, USB ports can be disabled for production use and then temporarily enabled to grant access to a crash cart for debugging purposes.

iDRAC Direct

iDRAC Direct is a special USB port that is hardwired to the iDRAC service processor for at-the-server debugging and management from the front of the server (cold aisle). It allows you to attach a standard Micro-AB USB cable to this port and the other end of the cable (Type A) to a laptop. A standard web browser can then access the iDRAC UI for extensive debugging and management of the server. If the iDRAC Enterprise license is installed, you can access the operating system desktop using the iDRAC's Virtual Console.

Because you use iDRAC credentials for logging in, iDRAC Direct works as a secure crash cart with the additional advantage of extensive hardware management and service diagnostics. This method is an attractive option for securing physical access to the server in remote locations (host USB ports and VGA outputs can be disabled in this case).

iDRAC Connection View with Geolocation

Connection View enables iDRAC to report the external switches and ports connected to Server I/O.

It is a feature on select networking devices and requires the Link Layer Discovery Protocol (LLDP) to be enabled on connected switches.

Some of the benefits of Connection View enable you to:

- Remotely and quickly check if server I/O modules (LOMs, NDCs, and add-in PCle cards) are connected to the correct switches and ports
- Avoid costly remote dispatch of technicians to remediate wiring errors
- Avoid tracing cables in the server room hot aisles
- Retrieve information for all connections by accessing the UI or by using RACADM commands


Beyond the obvious time and monetary savings, Connection View provides an additional benefit – real-time geolocation of a physical server or VM. Using iDRAC Connection View, administrators can pinpoint a server to see to which switch and port the server is connected. This information helps secure servers from being connected to networks and devices that do not comply with corporate security guidelines or best practices.

Connection View validates the location of the server indirectly by reporting the switch identities to which it is connected. The switch identity helps to determine the geolocation and to assure that the server is not a rogue server in a nonauthorized site, providing another layer of physical security. This information also provides validation that an application or VM has not "crossed" country borders, and that it is running in an approved, secure environment.

Protecting data at rest

Data at rest protection ensures that sensitive data that resides in storage is protected from unauthorized access through encryption and external key management.

Dell Technologies provides:

- Software-based encryption (for example, virtual devices)
- Enterprise key management (for example, SED devices and key management)
- Hardware drive encryption (for example, SED devices)

Whether it is due to internal policies or external compliance, securing data continues to be a high priority for organizations of all sizes.

The 14th, 15th, and 16th generations of PowerEdge servers offer several storage drive options for securing data, as shown in the following figure:





The options start with drives that support Instant Secure Erase (ISE), a new technology to erase user data instantly and securely. The 14th, 15th, and 16th generations of PowerEdge servers offer ISE-capable drives as a default. This white paper describes ISE in more detail later as part of the System Erase feature description.



The next higher security option is Self-Encrypting Drives (SEDs), which offer locking protection that binds the storage drive to the server and RAID card used. This method protects against so-called "smash and grab" theft of drives and the subsequent loss of sensitive user data. When thieves try to use the drive, they do not know the required locking key passphrase and are thwarted from accessing the encrypted drive data. Customers can protect against theft of the entire server by using Secured Enterprise Key Manager (SEKM), which is described in the following section.

NIST FIPS 140-2 certified SEDs offer the highest level of protection. Testing laboratories have accredited drives conforming to this standard. Tamper-resistant stickers are applied to the drive. Dell SED drives have FIPS 140-2 certification by default.

Secured Enterprise Key Manager

OpenManage Secured Enterprise Key Manager (SEKM) delivers a central key management solution to manage data at rest across the organization. It enables you to use an external Key Management Server (KMS) to manage keys that iDRAC can use to lock and unlock storage devices on a PowerEdge server. Using embedded code that is activated with a special license, iDRAC requests that the KMS creates a key for each storage controller, which iDRAC fetches and provides to the storage controller on every host boot so that the storage controller can unlock the SEDs.

The advantages of SEKM over Local Key Management (LKM) include:

- Protection against "Theft of a server" because the keys are not stored on the server and are stored externally and retrieved by connected PowerEdge server nodes (using iDRAC)
- Centralized and scalable key management for encrypted devices with high availability
- Support for the industry-standard Key Management Interoperability Protocol (KMIP), which enables the use of other KMIP-compatible devices
- · Protection of data at rest when drives or the entire server are compromised
- On-drive encryption performance scales with drive count



Figure 2. Secure Enterprise Key Manager (SEKM)

Local Key Management

PowerEdge servers provide the ability to secure SED drives connected to a PERC controller using Local Key Management (LKM).

To ensure user data protection if a drive is stolen, the SED must be locked with a separate key so that it does not decrypt user data unless that key is provided. This key is

referred to as the Key Encryption Key (KEK). The KEK is stored in the PERC, not on an external server.

Set a keyld/passphrase on the PERC controller to which the SED is connected. Then, the PERC controller generates a KEK using the passphrase and uses it to lock the SED. When the drive is powered on, it comes up as a locked SED and encrypts or decrypts user data only when the PERC provides the KEK to unlock it. If a locked drive is stolen, an attacker cannot provide the KEK, and the user data is protected.

The following figure shows the LKM solution:



Figure 3. Local Key Management (LKM)

iLKM

For direct attach NVMe configurations where a PERC RAID controller is not available, iDRAC can be used as the key manager. This solution called OpenManage iLKM, is iDRAC-based and enables key exchange locally. iDRAC acts as a key manager and generates authentication keys that can then be used to secure storage devices. You can transition from iDRAC-based iLKM to iDRAC-based SEKM to upgrade to external key management.

The following figure shows the iLKM solution:



Figure 4. Integrated Local Key Management (iLKM)

Protecting Data in Flight

Data-in-flight protection ensures that data is protected from unauthorized disclosure or interception as it travels across networks or between systems. Sensitive data can be intercepted, stolen, or modified in transit, leading to data breaches, loss of intellectual property, and other security risks.

Distributed and cloud environments where data is constantly moving between systems and across networks, protecting your data through encryption and access controls are important components for data-in-flight protection in a Zero Trust environment.



TLS 1.3

The iDRAC web server uses a TLS/SSL certificate to establish and maintain secure communications with remote clients. Web browsers and command-line utilities, such as RACADM and WS-Man, use this TLS/SSL certificate for server authentication and establishing an encrypted connection.

There are several options available to secure the network connection using a TLS/SSL certificate. iDRAC's web server has a self-signed TLS/SSL certificate by default. The self-signed certificate can be replaced with a custom certificate, a custom signing certificate, or a certificate signed by a well-known Certificate Authority (CA). Whichever method you choose, when iDRAC is configured and the TLS/SSL certificate is installed on the management stations, TLS/SSL-enabled clients can access iDRAC securely and without certificate warnings.

SSH

iDRAC provides user control over the cryptographic settings for the SSH daemon such that you can determine the ideal settings for your environment. The control given to you is not a relaxation of the settings. Instead, the feature enables you to modify the value set for each option to achieve a narrower and stringent cryptographic policy. That is, you can only remove values from the options but cannot add any values other than those values that have been defined and allowed in the default value-set.

The cryptographic policies are configured using the following options:

- Ciphers Ciphers
- Host-Key-Algorithms HostKeyAlgorithms
- Key-Exchange Algorithms KeyExchangeAlgorithms
- MACs MACs

Typically, the values for each of these options are set to prudent settings that reflect the best security practices that cater to a wide variety of environments. As such, the iDRAC default settings for these options are the same as those options assigned by the SSH package open-source community. These settings can be configured using the RACADM command-line interface. See the *iDRAC RACADM CLI User's Guide*.



Automatic certificate renewal

iDRAC9 v4.0 and later has added a client for Simple Certificate Enrollment Protocol (SCEP) support and requires a Datacenter License. SCEP is a protocol standard used for managing certificates for large numbers of network devices using an automatic enrollment process. iDRAC can now integrate with SCEP-compatible servers such as the Microsoft ServerNDES service to maintain SSL/TLS Certificates automatically. This feature can be used to enroll and refresh a soon-to-be-expired web server certificate. You can use Server Configuration Profile to set the certificates on a one-to-one basis in the iDRAC UI. Also, you can provide scripts using tools such as RACADM.





iDRAC Cipher Select

The Cipher Suite Selection can be used to limit the ciphers that the web browser can use to communicate with iDRAC. Also, it can determine the security of the connection. These settings can be configured through the iDRAC web interface, RACADM, and Redfish. This functionality is available across several iDRAC releases – iDRAC7, iDRAC8 (2.60.60.60 and higher), and the current iDRAC9 (3.30.30.30 and higher).

Commercial National Security Algorithm (CNSA) support

The supported ciphers available in iDRAC9 with TLS1.3 Bit and 256 Bit Encryption are shown in the following figure. The ciphers available are inclusive of the ciphers in the CNSA-approved set.

Support	ed Server	Cipher(s)		
Preferred	TLSv1.2	256 bits	ECDHE-RSA-AES256-GCM-SHA384	Curve P-256 DHE 256
Accepted	TL5v1.2	256 bits	ECDHE-RSA-AES256-SHA384	Curve P-256 DHE 256
Accepted	TLSv1.2	256 bits	ECDHE-RSA-AES256-SHA	Curve P-256 DHE 256
Accepted	TLSv1.2	256 bits	DHE-RSA-AES256-GCM-SHA384	DHE 2048 bits
Accepted	TL5v1.2	256 bits	DHE-RSA-AES256-SHA256	DHE 2048 bits
Accepted	TL5v1.2	256 bits	DHE-RSA-AES256-SHA	DHE 2848 bits
Accepted	TLSv1.2	256 bits	DHE-RSA-CAMELLIA256-SHA	DHE 2048 bits

Tips for securing iDRAC connection

The most secure network connection is the iDRAC's Dedicated NIC because it can be connected to a network that is physically separated from the production network. This method physically separates the iDRAC management traffic from the production network traffic.

If use of the iDRAC's Dedicated NIC is not feasible, iDRAC can be run in Shared LOM mode with a VLAN enabled. However, the iDRAC's management traffic is sent across the same connection as the production network. Alternatively, if the use of a VLAN is not possible while in Shared LOM mode, access to iDRAC must be secured using strong passwords and other security measures.

For further information, see the iDRAC Security Configuration Guide.

IEEE.802.1x

The network standard IEEE 802.1x has been enabled on PowerEdge servers. This network protocol provides port-based network authentication. Devices requesting access to the LAN or WLAN must be authenticated and validated before authorization for access.

Domain isolation

The 14th, 15th, and 16th generations of PowerEdge servers provide additional security by using domain isolation, an important feature for multitenant hosting environments. To secure the server's hardware configuration, hosting providers might want to block any reconfiguration by tenants. Domain isolation is a configuration option that ensures that management applications in the host operating system have no access to the out-of-band iDRAC or to Intel chipset functions, such as Management Engine (ME) or Innovation Engine (IE).

Protecting data in use

Protecting application data that is being used in memory has become increasingly important. Whether data in use is a machine learning dataset or relates to keeping a secret in memory such as in multitenant environments, data in-use can be vulnerable to

threat vectors that can intrude on the contents of memory or the access bus. Data in use protection is necessary to secure computations that are increasingly operating on large datasets in memory. Additionally, code running on the data must be trusted and tamper-free. There must be facilities to separate trusted and nontrusted code execution environments for data in use.

New developments in CPU technology for confidential computing allow secure enclaves to protect application data at the hardware layer, enabling a more comprehensive data protection strategy. Starting with the 15th generation of server offerings, Dell Technologies enables these CPU technologies, including Intel SGX/TME and AMD SEV/SME.

AMD confidential compute features

AMD introduced Secure Encrypted Virtualization (SEV) with the first generation of AMD EPYC processors. It encrypts full system memory and individual virtual machine (VM) memory isolating the VM memory from the hypervisor. With each generation of AMD EPYC processors, AMD has enhanced SEV with additional features to safeguard privacy and integrity by encrypting each VM with one of up to 509 unique encryption keys known only to the processor. The keys are used to encrypt memory using 128-bit AES encryption engines in the memory controller. The hypervisor manages the keys in the memory controller with the help of AMD's Secure processor. The key AMD SEV technologies that are part of the Infinity Guard technology solution suite from AMD that address different use cases, deployments, and threat models include:

- AMD's SME technology refers to using a single key to encrypt system memory.
- AMD SEV uses one key per VM to isolate guests and the hypervisor from one another.
- AMD SEV-ES encrypts all CPU register contents when a VM stops running to prevent leakage of information from CPU registers to the hypervisor.
- AMD SEV-SNP adds strong memory integrity protection to help prevent hypervisor-based attacks like data replay, memory remapping, and so on.

Intel confidential computing features

Starting with the 3rd Generation Intel Xeon platform, Intel introduced several key security innovations. Total Memory Encryption (TME) is available to ensure memory accessed from the CPU is encrypted. By encrypting all memory, existing software applications run unmodified while simultaneously providing greater protection for system memory.

Intel TME helps ensure that all memory accessed from the Intel CPU is encrypted, including customer credentials, encryption keys, and other intellectual property or personal information about the external memory bus. Intel developed this feature to provide greater protection for system memory against hardware attacks, such as removing and reading the dual inline memory module (DIMM) after spraying it with liquid nitrogen or installing purpose-built attack hardware. Using the NIST storage encryption standard, AES XTS, an encryption key is generated using a hardened random number generator in the processor without exposure to software. This method allows existing software to run unmodified while better protecting memory.



Identity Access Management

Identity and Access Management (IAM) is a set of security controls to manage digital identities and provide authentication and authorization by controlling the requestor's access to information and resources, at the right level, and at the right time to limit unauthorized access. IAM frameworks provide advantages over simple passwords, such as:

- Enhanced security—Includes various tools such as single sign-on (SSO) services, multifactor authentication (MFA), and privileged access management for stronger security and to avoid risks from vulnerable passwords. IAM plays a critical role in protecting organizations against phishing and social engineering attacks by implementing strong authentication, access control, monitoring, and user education. It helps organizations reduce their attack surface and respond effectively to threats.
- **Reduction of compromised passwords**—Deters credential theft from phishing, social engineering, and brute-force attacks by using MFA and additional security layers.
- **Granular access control** Enables fine-grained control over user access permissions by using role-based access control (RBAC) and attribute-based access control (ABAC) to ensure that users only have access to necessary resources and data.
- **Centralized management**—Controls user accounts and access policies, which provides easier management of users as they join, move within, or leave the organization.
- Auditing and logging—Monitors user activities to identify and remediate suspicious access attempts.
- Scalability-Adapts to the need for a growing number of users and resources
- Regulatory compliance—Adheres to regulatory requirements to avoid legal consequences.
- User experience—Provides SSO services and password management to enable users to access multiple applications and services by using one set of credentials.
- Adaptive authentication—Applies additional security measures by assessing risk factors.
- Emergency access and recovery—Grants temporary emergency access for critical situations while maintain security controls.
- Integration—Integrates with various systems, applications, and cloud services for seamless access management across the entire IT ecosystem.

As a broad IT issue spanning technological and regulatory requirements, IAM is a strategic business imperative for all organizations to enhance security and resiliency. Zero Trust Architecture (ZTA) has emerged as the standard choice for securing all levels of infrastructure and is a foundational part of ZTA. The unofficial mantra of Zero Trust is "never trust, always verify" and if you cannot verify then you cannot trust. IAM provides that verification.



The foundation to a strong cybersecurity framework and the adoption and implementation of Zero Trust principles is identity. Identification of not only people but of applications, communication paths, network devices, and the data itself. If an organization does not have a strong Identity Credentialing and Access Management (ICAM) practice, then the underlying security practices are at risk. An effective identity and access solution must include necessary tools and controls that can capture and store user login details, facilitate the assignment and revocation of user access credentials, and oversee the central enterprise database of user roles, levels, and access privileges. At a minimum, access must only be allowed to facilitate the necessary functions that the organization has defined. As an example, users must only have access and permission to data, applications, and services to do what is defined in their job. Practices and tools must be implemented to ensure that this role is maintained and that even in that role, a user must be monitored to ensure that they are not doing anything that violates the organization's goals for data utilization.

MFA – Smartcards (CAC/PIV)

MFA for Smartcards (CAC/PIV) is a general-purpose certificate authentication that includes Common Access Card (CAC) and Personal Identification Verification (PIV) cards. Certificate authentication uses the client identity certificate to authenticate the user. It is used primarily in government or organizations that work with the federal industry.

MFA – RSASecurID

RSA SecurID is another means of authenticating a user on a system. As another twofactor authentication, iDRAC9 supports RSA SecurID with the Datacenter license and starting with firmware 4.40.00.00 and later, as another two-factor authentication.

Directory integration for authorization

For centralized user and domain management, iDRAC supports integration into privileged management tools such as Lightweight Directory Access Protocol (LDAP) and Active Directory. Using a directory service provides a central location for easier user inventory management and assigning user account access controls and settings.

You can use LDAP to authenticate users and groups in iDRAC. To configure the LDAP directory service, you can use the objects in the cfgLdap and cfgLdapRoleGroup groups with the config command. You can also use the objects in the iDRAC.LDAP and iDRAC.LDAPRole groups with the set command 12.

For Active Directory integration, you can configure LDAP over SSL (LDAPS) on iDRAC to communicate with Active Directory domain controllers. After a valid certificate is installed on the domain controller and the connection to the DC using SSL over port 636 is verified, you can use the directory service integration test on iDRAC/OME to communicate with the domain controller.

SSO

iDRAC supports SSO, which provides the ability to share validated credentials and identifications across multiple domains without having to rechallenge or reauthenticate the user. SSO enables an authenticated operating system administrator to directly access the iDRAC web interface without requiring login using separate iDRAC administrator credentials. iDRAC supports the following SSO protocols:

• **OpenID connect** is an open standard and is a decentralized authentication protocol that is typically used for machine-to-machine-like RestAPIs.



• Open-standards **Security Assertion Markup Language (SAML)** is typically used for UI SSO.

Role-based access controls

Role-based access control (RBAC) is the most popularly used form of access control. Permissions are grouped in roles and are typically assigned to a group. Assigning authorization capabilities to users and groups is managed in Active Directory.

You can set up user accounts with specific privileges (role-based authority to manage your system using iDRAC and maintain system security). By default, iDRAC is configured with a local administrator account. The default iDRAC username and password are provided with the system badge. As an administrator, you can set up user accounts to allow other users to access iDRAC. For more information, see the documentation for the server.

You can set up local users or use directory services such as Microsoft Active Directory or LDAP to set up user accounts. Using a directory service provides a central location for managing authorized user accounts.

iDRAC supports role-based access to users with a set of associated privileges. The roles are administrator, operator, read-only, or none. The role defines the maximum privileges available.

Time-based access controls

Time-based access control is another valuable tool for enhancing security and managing access to sensitive data, facilities, or systems in a tightly controlled and monitored environment. For instance, if a technician needs physical access to the USB port on the server, the iDRAC administrator can enable/disable specific access times for USB port access.

Scope-based access control

Scope-based access control provides more granular control than user- and role-based access control. It allows the administrator to apply a policy with a set of permissions that are evaluated when an entity tries to access the resource. For instance, access can be restricted, based on resource location, IP address range, and so on.

Capabilities and automation for efficient at-scale deployment

Zero-touch automation with iDRAC – Server configuration profiles

Zero Touch provisioning is available with iDRAC Enterprise or Datacenter licenses. Zero Touch provisioning automates all hardware configuration, certificate installation, repository firmware updates, and operating system deployment. The IT Admin can preconfigure security settings and ensure uniform server images. Zero Touch provisioning is available through the iDRAC Server Configuration Profile feature and with OpenManage Enterprise.

CloudIQ

CloudIQ monitors the health and cybersecurity of your enterprise-wide servers and predicts their performance so that you can proactively address issues before they impact your business.

CloudIQ offers a simple and intuitive solution to collect firmware details from PowerEdge servers including BIOS, iDRAC, NICs, PERC, drives, and supported peripherals. A recent

30 Cyber Resilient Security in Dell PowerEdge Servers

Inserido ao protocolo **22.951.206-4** por: **Pedro Henrique Golin Linhares** em: 31/03/2025 12:49. A autenticidade deste documento pode ser validada no endereço: https://www.eprotocolo.pr.gov.br/spiweb/validarDocumento com o código: 692a5ef82beb173ed33dad4ee6505546.

feature from both OpenManage Enterprise and CloudIQ identifies BIOS and firmware that requires an update. CloudIQ can report the current installed version, compare it to the latest Dell release available, and schedule updates. This information is collected from each server using the agent-free iDRAC, consolidated by OpenManage Enterprise, and then transferred to CloudIQ to be processed. This powerful feature includes user rights integrated in both OpenManage Enterprise and CloudIQ to allow only authorized users to run these commands. Also, CloudIQ can consolidate multiple OpenManage Enterprise instances into one server fleet management view.

Fourth Stage – Security management and monitoring

Challenges

You cannot defend what you cannot see. Attacks happen quickly, faster than a human can detect and respond. Protecting your critical assets in real-time within a dynamic and complex threat environment is a challenge. Lack of skills and resources exacerbates the problem. As IT administrators face ongoing challenges in their environments, more easily managing the infrastructure—through more automation, fewer task steps, and more intuitive interactions—is key to administration productivity.

It is critical that your business remains resilient and unaffected by adverse outcomes. The business must stay nimble to be aware, respond, and recover:

- Observability and transparency enable effective security event awareness and timely remediation.
- Lack of skills and resources increases problems. Monitor activity across infrastructure (user, device, data, network, and application) to identify patterns and anomalies. Analytics can be used to detect and respond to security threats.
- Automation is important.
- Analyze events, activities, and behaviors to derive context and apply Al/ML to achieve models that improve detection and reaction time in making real-time access decisions. Simplification is foundational and fundamental to a resilient infrastructure.

PowerEdge Solutions

The Dell management portfolio simplifies administrator tasks. It improves security and health monitoring to scale security confidently by using automation and intelligence. Monitoring and logging are a key part of a zero-trust implementation.

Visibility, logging, and alerts

It is critical to have a detection capability that provides complete visibility into the configuration, health status, and change events in a server system. This visibility must also detect malicious or other changes to BIOS, firmware, and option ROMs in the boot and operating system runtime process. Proactive polling must be coupled with the ability to send alerts for any events in the system. Logs must provide complete information about access and changes to the server. Most importantly, the server must extend these capabilities to all components.



Dell OpenManage Enterprise enables users to set up alert policies once and then automatically assign them for future alerts. Also, OpenManage Enterprise can apply a template to many servers at once. The OpenManage Enterprise solution ultimately saves time and effort by automating actions based on alerts after administrators have created a policy.

Telemetry

Beginning with iDRAC9 v4.00.00.00 firmware and a Datacenter license, IT managers can integrate advanced server hardware operation telemetry into their existing analytics solutions. Telemetry is provided as granular, timeseries data that is streamed or pushed. The advanced agent-free architecture in iDRAC9 provides over 180 data metrics that are related to server and peripheral operations. Metrics are precisely timestamped and internally buffered to allow highly efficient data stream collection and processing with minimal network loading. This comprehensive telemetry can be fed into analytics tools to predict failure events, optimize server operations, and enhance cyber resiliency. The iDRAC9 Telemetry Streaming collects and streams live system data from one or more PowerEdge servers to a centralized collector.

iDRAC Lifecycle Logs

The Lifecycle Logs are a collection of events that occur in a server over time. They provide a description of events with timestamps, severity, user ID or source, and recommended actions. This technical information helps with security tracking and other hardware alerts.

The various types of information that is recorded in the Lifecycle Controller Log (LCL) include:

- Configuration changes on the system hardware components
- iDRAC, BIOS, NIC, and RAID configuration changes
- Logs of all the remote operations
- Firmware update history based on device, version, and date
- Information about replaced parts
- Information about failed parts
- Event and error message IDs
- Host power-related events
- POST errors
- User login events
- Sensor state change events

Alerts

iDRAC provides the capability to configure different event alerts and actions to be performed when a Lifecycle Logs event occurs. When an event is generated, it is forwarded to the configured destinations by using the selected alert type mechanisms. Users can enable or disable alerts through the iDRAC web interface, RACADM, or with the iDRAC settings utility.



iDRAC supports several types of alerts such as:

- Email or IPMI alert
- SNMP trap
- Operating system and Remote System logs
- Redfish event

Alerts are categorized by severity – Critical, Warning, or Informational. The following filters can be applied to alerts:

- System health-For example, temperature, voltage, or device errors
- Storage health-For example, controller errors, physical or virtual disk errors
- **Configuration changes**—For example, change in RAID configuration, PCIe card removal
- Audit logs-For example, password authentication failure
- **Firmware**—For example, upgrades or downgrade

The IT administrator can set different actions for alerts – Reboot, Power Cycle, Power Off, or No action.

TLS for Remote Syslog

The iDRAC Remote Syslog feature allows you to write the RAC log and the System Event Log (SEL) remotely to an external syslog server. You can read all logs from the entire server farm from a central log. The Remote Syslog protocol does not require user authentication. For the logs to be entered in the Remote Syslog server, ensure that there is proper network connectivity between iDRAC and the Remote Syslog server and that the Remote Syslog server is running on the same network as iDRAC.

The iDRAC's web server has a self-signed TLS/SSL certificate by default. The selfsigned certificate can be replaced with a custom certificate, a custom signing certificate, or a certificate signed by a well-known Certificate Authority (CA). Redfish scripts can perform automated certificate uploads. When a link has been established between the two servers, TLS encryption and SSL decryption enable secure data transport.

SELinux framework

SELinux operates at the core kernel level on iDRAC and does not require user input or configuration. SELinux logs security messages when an attack is detected. These log messages indicate when and how an attacker tried to break into the system. These logs are available through SupportAssist to customers enrolled in this new feature. In future releases of iDRAC, these logs will be available in the Lifecycle Controller Logs.

Real-time detection – BIOS Live scan

BIOS live scanning verifies the integrity and authenticity of the BIOS image in the primary ROM when the host is powered on. BIOS live scanning is not in the POST process. This feature is available only with the iDRAC9 4.10.10.10 (supported AMD platforms) and iDRAC9 4.40.20.00 (supported Intel platforms) Datacenter licenses. You must have administrator privileges or operator privileges with the "Execute Debug Commands" debug privilege to perform this operation.



You can initiate BIOS image scanning either on demand or by scheduling the scan through the iDRAC UI, RACADM, and Redfish interfaces. The BIOS live scan feature is available starting with 15th generation of PowerEdge servers with AMD "Rome"-based processors or Intel "Ice Lake"-based processors.

Boot time and run time BIOS scanning

A critical aspect of server security is ensuring that the boot process is verified as secure. This process provides a trusted anchor for all subsequent operations such as booting an operating system or updating firmware. PowerEdge servers have used silicon-based security for several generations for features such as iDRAC Credential Vault, an encrypted secure memory in iDRAC for storing sensitive data. The boot process is verified using a silicon-based RoT to meet the following recommendations:

- NIST SP 800-147B, BIOS Protection Guidelines for Servers
- NIST SP 800-155, BIOS Integrity Measurement Guidelines

On the Dell PowerEdge servers with iDRAC9, iDRAC first boots with chain of trust authentication, and then verifies BIOS integrity. iDRAC assumes the role of hardwarebased root of trust. For AMD platforms, iDRAC accesses the primary BIOS ROM through SPI and the AMD fusion controller hub (FCH), and performs the RoT process. For Intel platforms, iDRAC accesses the primary BIOS ROM through SPI and the Intel Platform Controller Hub (PCH), and performs the RoT process.

iDRAC9 directly accesses the BIOS primary ROM to perform a RoT operation on the processor on both the security block and the host Initial Boot Block.

Silicon-based Root of Trust

PowerEdge servers use an immutable, silicon-based RoT to attest to the integrity of BIOS and iDRAC9 firmware cryptographically. This RoT is based on one-time programmable, read-only public keys that provide protection against malware tampering. The BIOS boot process uses Intel Boot Guard technology or AMD Platform Secure Boot technology. This technology verifies that the digital signature of the cryptographic hash of the boot image matches the signature stored in silicon by Dell Technologies in the factory. A verification failure results in a server shutdown and user notification in the Lifecycle Controller Log. The user can initiate the BIOS recovery process. If Boot Guard validates successfully, the other BIOS modules are validated by using a chain of trust procedure. Then, control is given to the operating system or hypervisor. In addition to Boot Guard, iDRAC9 4.10.10.10 or later provides a RoT mechanism that verifies the BIOS image at the host boot time. The host can boot only after the BIOS image is successfully validated. iDRAC9 also provides a mechanism to validate the BIOS image at run time, on demand, or at user-scheduled intervals.

For the chain of trust, each BIOS module contains a hash of the next module in the chain. The key modules in the BIOS include:

- Technical support and resources ID 483
- Initial Boot Block (IBB)
- Security (SEC)
- Pre-EFI Initialization (PEI)
- Memory Reference Code (MRC) o Driver Execution Environment (DXE)

34 Cyber Resilient Security in Dell PowerEdge Servers

Inserido ao protocolo **22.951.206-4** por: **Pedro Henrique Golin Linhares** em: 31/03/2025 12:49. A autenticidade deste documento pode ser validada no endereço: https://www.eprotocolo.pr.gov.br/spiweb/validarDocumento com o código: 692a5ef82beb173ed33dad4ee6505546.



• Boot Device Selection (BDS)

If Intel Boot Guard authenticates the IBB module, the IBB module validates the SEC and PEI modules before handing control to it. The SEC and PEI modules then validate the PEI and MRC modules, which further validates the DXE and BDS modules. Next, control is handed over to UEFI Secure Boot. Similarly, for PowerEdge AMD EPYC-based servers, AMD Secure Root of Trust technology ensures that servers boot only from trusted firmware images. AMD Secure Run Technology encrypts the main memory, keeping it private from malicious intruders accessing the hardware. There are no required application modifications to use this feature, and the security processor never exposes the encryption keys outside of the processor. iDRAC takes on the role of hardware-based security technology and accesses the primary BIOS ROM through SPI. iDRAC, along with the AMD fusion controller hub (FCH) performs the RoT process.

Under the following conditions, iDRAC9 recovers the BIOS:

- BIOS integrity check failed
- BIOS self-check failed

Note: Use the RACADM command to recover the BIOS setup.

The iDRAC boot process uses its own independent silicon-based RoT that verifies the iDRAC firmware image. The iDRAC RoT also provides a critical trust anchor for authenticating the signatures of Dell firmware update packages (DUPs).

System lockdown

iDRAC9 offers a feature that 'locks down' the server hardware and firmware configuration and requires an Enterprise or Datacenter license. You can enable this mode by using the UI, the RACADM CLI, or the Server Configuration Profile. Users with administrative privileges can set System Lockdown mode, which prevents users with lesser privileges from changing the server. The IT administrator can enable or disable this feature. Any changes made when System Lockdown is disabled are tracked in the Lifecycle Controller Log. By enabling lockdown mode, you can prevent configuration drift in your data center when using Dell tools and agents and protect against malicious attacks against embedded firmware when using Dell Update Packages. Lockdown mode can be enabled dynamically, without requiring a system reboot. iDRAC9 v4.40 introduced enhancements where in addition to the current System Lockdown which only controls the updates using Dell Update Package (DUP), this lockdown functionality is extended to select NICs.

Note: Enhanced Lockdown for NICs only includes firmware lockdown to prevent firmware updates.

Configuration (x-UEFI) lockdown is not supported. When the customer sets the system in lockdown mode by enabling or setting attributes from any of the supported interfaces, iDRAC will take additional actions depending on the system configuration. These actions depend on the third-party devices detected as part of the iDRAC discovery process.

Drift detection

By enforcing standardized configurations and adopting a "zero tolerance" policy for any changes, organizations can reduce the potential for exploitation. Dell OpenManage



Enterprise Console allows you to define your own server configuration baselines and then monitor the drift of your production servers from those baselines. The baseline can be built based on different criteria to fit different production enforcement, such as security and performance.

OpenManage Enterprise can report any deviations from the baseline and optionally repair the drift with a simple workflow to stage the changes on iDRAC out of band. The changes can then take place at the next maintenance window while servers reboot to make the production environment compliant again. This staged process enables you to deploy configuration changes to production without any server downtime during nonmaintenance hours. It increases the server availability without compromising the serviceability or security.

Chassis intrusion detection

PowerEdge servers provide hardware intrusion detection and logging, with detection working even when no AC power is available. Sensors on the chassis detect when anyone opens or tampers with the chassis, even during transit. Servers that have been opened while in transit generate an entry in the iDRAC Lifecycle Logs after power is supplied.

Automated and manual recovery

BIOS and operating system recovery

The 14th, 15th, and 16th generations of PowerEdge servers include two types of recovery: BIOS Recovery and Rapid Operating System (OS) Recovery. These features enable rapid recovery from corrupted BIOS or operating system images. In both cases, a special storage area is hidden from run-time software (BIOS, operating system, device firmware, and so on). These storage areas contain pristine images that can be used as alternatives to the compromised primary software.

Rapid Operating System (OS) Recovery enables rapid recovery from a corrupted operating system image (or an operating system image suspected of malicious tampering). The recovery media can be accessed using an internal SD card, SATA ports, M.2 drives, or internal USB. The selected device can be exposed to the boot list and the operating system to install the recovery image. It can then be disabled and hidden from the boot list and operating system. In the hidden state, the BIOS disables the device so that the operating system cannot access it. If there is a corrupted operating system image, the recovery location can then be enabled for the boot process. These settings can be accessed through BIOS or the iDRAC interface.

In extreme cases, if the BIOS is corrupted (either by a malicious attack, a power loss during the update process, or any other unforeseen event), it is important to provide a way to recover the BIOS to its original state. A backup BIOS image is stored in iDRAC so it can be used to recover the BIOS image if needed. iDRAC orchestrates the entire end-to-end recovery process.

- BIOS itself initiates automatic BIOS recovery.
- Users using the RACADM CLI command can initiate on-demand BIOS recovery.

FW rollback

We recommend that you update the firmware to ensure you have the latest features and security updates. However, you may need to rollback an update or install an earlier

version if you encounter issues after an update. If you roll back to the previous version, it is also verified against its signature.

Firmware Rollback from existing production version "N" to a previous version "N-1" is supported for the following firmware images:

- BIOS
- iDRAC with Lifecycle Controller
- Network Interface Card (NIC)
- PowerEdge RAID Controller (PERC)
- Power Supply Unit (PSU)
- Backplane

You can roll back the firmware to the previously installed version (N-1) using any of the following methods:

- iDRAC web interface
- CMC web interface
- RACADM CLI for iDRAC and CMC
- Lifecycle Controller UI
- Lifecycle Controller remote services

You can roll back the firmware for iDRAC or any device that Lifecycle Controller supports, even if the upgrade was previously performed using another interface. For example, if the firmware was upgraded using the Lifecycle Controller UI, you can roll back the firmware using the iDRAC web interface. You can perform firmware rollback for multiple devices with one system reboot.

On 14th, 15th, and 16th generations of PowerEdge servers that have a single iDRAC and Lifecycle Controller firmware, rolling back the iDRAC firmware also rolls back the Lifecycle Controller firmware.

Firmware rollback protection

If the firmware has a known vulnerability which would expose your server to attack, the BIOS itself can prevent downgrade to a previous version. The firmware release notes state that you cannot rollback when performing an update.

Full Power Cycle

In a Full Power Cycle, the server and all its components are rebooted. It drains main and auxiliary power from the server and all components. All data in volatile memory is also erased.

A physical Full Power Cycle requires removing the AC power cable, waiting for 30 seconds, and then putting the cable back. This method poses a challenge when working with a remote system. A feature in the 14th, 15th, and 16th generations of PowerEdge servers enables you to perform an effective full power cycle from the iDRAC Service Module (iSM), IDRAC UI, BIOS, or a script. A full power cycle takes effect at the next power cycle.



The Full Power Cycle feature eliminates the need for anyone to be physically present in the data center, thus reducing time to troubleshoot. It can eliminate, for example, any malware that is still memory-resident.

Updating

Dell Technologies provides a rich set of tools to make it easier to keep your server's firmware and BIOS up-to-date and update quickly. Ensuring firmware is up to date is a vital task to keep production servers secure and operating efficiently. Tracking and implementing these updates can be burdensome for administrators. iDRAC9 provides automatic updates with the ability to schedule firmware updates as wanted.

Many companies schedule monthly maintenance windows to handle operating system, application, and firmware updates. With OpenManage Enterprise, systems administrators can stage firmware updates for the next time the system is rebooted or for a scheduled deployment. This method ensures that nobody must be physically present to run the updates.

PowerEdge provides patches to security vulnerabilities with Dell security advisories. The advisories provide timely information, guidance, and mitigation options to minimize risks associated with security vulnerabilities.

Restoring server configuration after hardware servicing

Remediating service events is a critical part of any IT operation. The ability to meet recovery time objectives and recovery point objectives has direct implications on the security of the solution. Restoring server configuration and firmware assures that security policies for server operation are automatically met.

PowerEdge servers provide functionality that quickly restores server configuration in the following situations:

- Individual part replacement
- Motherboard replacement (full server profile backup and restore)
- Motherboard replacement (easy restore)

Parts replacement

iDRAC automatically saves the firmware image and configuration settings for NIC cards, RAID controllers, and Power Supply Units (PSUs). If there is a field replacement of these parts, iDRAC automatically detects the new card and restores the firmware and configuration to the replaced card. This functionality saves critical time and ensures a consistent configuration and security policy. The update occurs automatically on system reboot after replacing the supported part.

Easy Restore (for motherboard replacement)

Easy Restore is an integrated storage component that maintains critical configuration information. The Easy Restore feature allows you to restore your system's service tag, all licenses, UEFI configuration, system configuration settings (BIOS, iDRAC and NIC), and the OEM ID (Personality Module) after replacing the system board. All data is backed up to a backup flash device automatically. If the BIOS detects a new system board and the service tag in the backup flash device, the BIOS prompts you to restore the backup information. You can still choose to perform the full system backup with iDRAC8 as you do with iDRAC6 and iDRAC7. This solution backs up and restores the actual firmware

versions in addition to the hardware settings. Easy Restore does not copy the firmware drivers due to size limitations.

CloudIQ

Misconfigurations of infrastructure systems can open your organization to cyber intrusion and are a leading threat to data security. The CloudIQ cybersecurity feature proactively monitors infrastructure security configurations for Dell PowerStore and PowerMax storage systems and PowerEdge servers, and notifies users of security risks. A risk level is assigned to each system, placing the system into one of four categories, depending on the number and severity of the issues: Normal, Low, Medium, or High.

By using CloudIQ Cybersecurity policy templates, users can quickly set up security configuration evaluation tests and assign them to large numbers of systems with just a few clicks. Once assigned, the test plan is evaluated against each associated system, and the system administrator is notified in minutes of any unwanted configuration settings.

When a security risk is found, remediation instructions are provided to help you address the issue quickly. CloudIQ evaluates outgoing Dell Security Advisories (DSAs) and intelligently notifies users when those advisories are applicable to their specific Dell system models with specific system software and firmware versions. This notification eliminates the need for users to investigate if a Security Advisory applies to their systems and allows them to focus on remediation immediately.

Managed detection and response services

Dell Technologies Managed Detection and Response services is a cloud-based offering that helps organizations quickly and significantly improve their security posture—while reducing the burden on IT. This fully managed, end-to-end, 24/7 service monitors, detects, investigates, and responds to threats across the entire IT environment.

Designed for organizations with 50 endpoints or more, this unique service uses two key capabilities:

- The power of the open Secureworks Taegis XDR security analytics software, built on more than 20 years of SecOps expertise including real-world threat intelligence and research, and experience detecting and responding to advanced threats
- The expertise of Dell Technologies security analysts, gained through years of experience helping organizations worldwide to better protect their business

Dell Technologies Managed Detection and Response provides around-the-clock access to security experts. They provide end-to-end visibility and protection from endpoint to cloud, covering every aspect of advanced threat detection and response supported by Taegis XDR's database of 52,000 unique threat indicators that are managed and updated daily. Taegis XDR also ingests data from existing security solutions to use any previous security investments.

Dell Technologies security analysts assist with initial setup, monitoring, detection, remediation, and response—all for one predictable price. They work closely with your IT team to understand the environment. They provide advice about improvements to the security posture and help you set up and deploy the Taegis XDR software agent to endpoints. Then, using the Taegis XDR application, they monitor and review alerts 24/7.



Fifth stage - Secure decommissioning and repurposing

If an alert merits investigation, analysts determine and perform the appropriate response. If a threat is malicious or requires your action, you are informed and, if necessary, provided with step-by-step instructions. As part of the service, Dell Technologies also provides up to 40 hours per quarter of <u>remote remediation assistance</u>, such as helping with troubleshooting, issue resolution, software deployments, patch and asset assessment, and configuration of IT environments.

If there is a security incident, Dell Technologies initiates the process to get your business up and running and provides up to 40 hours of remote incident response assistance a year.

Fifth stage – Secure decommissioning and repurposing

Challenge Data security is a key consideration throughout the life cycle of a server, including when the server is repurposed or retired. Many servers are repurposed as they are transitioned from workload to workload, or as they change ownership from one organization to another. All servers are retired when they reach the end of their useful life. When such transitions occur, IT best practices recommend removing all data from the server to ensure that confidential information is not inadvertently shared.

PowerEdge solutions

Beyond best practices, in many cases government regulations about privacy rights also necessitate complete data elimination when IT resources are transitioned. Data erasure is a key capability encompassed in the Dell Secure Development Lifecycle (SDL). The SDL and secure server management tools ensure that PowerEdge servers are secure at every stage in the server life cycle, from server conception, design and manufacturing, to operation and decommissioning.

At this final stage (decommissioning/retirement), or when a server is repurposed due to change of workload or ownership, a capability starting with the PowerEdge 14th generation of servers, can simplify data erase.

System Erase, with iDRAC9 and the 14th, 15th, and 16th generations of PowerEdge servers, simplifies the process of erasing server storage devices and server nonvolatile stores such as caches and logs. To meet varying Systems Administrator needs for interactive and programmable operations, the following methods can perform System Erase:

- Lifecycle Controller UI
- WS-Man API
- RACADM CLI

Using one of these methods, an administrator can selectively reset a PowerEdge server to its original state (factory settings), removing data from internal server nonvolatile stores and from storage devices within the server.

System Erase can discover server-attached storage including hard disk drives (HDDs), SEDs, ISE, and nonvolatile memory drives (NVMes). Data stored on ISE, SED, and NVMe devices can be made inaccessible using cryptographic erase while devices such as non-ISE SATA HDDs can be erased using data overwrite.

40 Cyber Resilient Security in Dell PowerEdge Servers

Inserido ao protocolo **22.951.206-4** por: **Pedro Henrique Golin Linhares** em: 31/03/2025 12:49. A autenticidade deste documento pode ser validada no endereço: https://www.eprotocolo.pr.gov.br/spiweb/validarDocumento com o código: 692a5ef82beb173ed33dad4ee6505546.



Secure Erase

Through the life cycle controller, customers can repurpose or retire a system. All drives now shipping on PowerEdge systems can be securely erased. On an older platform that might not have encryption-capable drives, there is a "standard disks (overwrite data)" option. Unrecoverable processes generate warning messages. The server is powered off after a retire or repurpose operation. You can view the Lifecycle Logs in iDRAC to confirm that the operation was successful.

At the end of a system life cycle, it can either be retired or repurposed. For either scenario, System Erase removes sensitive data and settings from the server. Secure Erase wipes storage devices and server nonvolatile stores such as caches and logs so that no confidential information unintentionally leaks. It is a utility in Lifecycle Controller (F10) that erases logs, configuration data, storage data, and cache.

The System Erase feature can erase the following devices, configuration settings, and applications:

- iDRAC is reset to default settings, erasing all data and settings.
- Lifecycle Controller (F10) data is cleared.
- BIOS and NVRAM are reset to default settings.
- Embedded diagnostics and operating system driver packs are cleared.
- iDRAC Service Module (iSM) are cleared.
- SupportAssist Collection reports are cleared.

The following components can also be erased:

- Hardware Cache (clear PERC NVCache)
- vFlash SD Card (initialize card)

Note: vFlash not available on 15th generation of PowerEdge servers or later.

System Erase cryptographically disposes of data on the following components:

- Self-Encrypting Drives (SED)
- ISE drives
- NVM devices such as Intel Apache Pass and NVDIMMs

Secure erase – physical disk

Reset the drive to the factory settings. All data on the SSD is permanently removed and cannot be recovered. To sanitize the drive, the mapping table is deleted and all blocks to which data has been written are erased. Not all SSDs support the sanitize feature.

Overwrite erase

Overwrite-erase is a software-based method that overwrites data with zeros and ones. Data overwrite can erase non-ISE SATA hard drives.



Summary

Crypto erase

ISE destroys the internal encryption key that is used in the 14th, 15th, and 16th generations of PowerEdge drives and renders the user data unrecoverable. Used on self-encrypting drives, the encryption key is erased. The data remains on the drive but is inaccessible without the key. ISE is a recognized method of data erasure on storage drives as seen in NIST Special Publication 800-88 "Guidelines for Media Sanitization."

Advantages of the new ISE feature with System Erase include:

- **Speed**—ISE is faster than data overwriting techniques like DoD 5220.22-M (seconds compare with hours)
- **Effectiveness**—ISE renders all the data on the drive, including reserved blocks, unreadable.
- Better TCO—Storage devices can be reused instead of being physically destroyed.

The following methods can perform System Erase procedure:

- Lifecycle Controller interface (F10)
- RACADM CLI
- Redfish

Data sanitization and destruction services

Data continues to grow and drive strategic advantage. Meanwhile, national security issues and data privacy regulations are escalating. As organizations navigate technology changes, they are further challenged with data security and compliance.

The end of the product life cycle is an aspect of data security that is increasingly important. Data Sanitization for Enterprise is a software-based method of securely overwriting data to render it unrecoverable. The various options include:

- Data Sanitization for Enterprise Onsite—An option for customers looking to refresh or redeploy assets. This service performs sanitization at the business' location, securing data while assets remain in the environment.
- Data Sanitization for Enterprise Offsite with Asset Resale and Recycle—A service that removes assets from the business' environment, sanitizes them at a secure location, and evaluates them for resale or reuse. The customer is compensated if value is found. If no value is found, assets are recycled according to local regulatory guidelines.

Data Destruction for Enterprise—A process that renders data inaccessible through the process of physical shredding. It is available for all Dell infrastructure solutions and similar third-party non-Dell branded assets. This process does not require systems to be operational.

Summary

Data center security is paramount to business success, and the security of the underlying server infrastructure is critical. Cyberattacks have the potential for extended system and business downtime, lost revenue and customers, legal damages and tarnished corporate

reputation. To protect, detect, and recover from hardware-targeted cyberattacks, security must be built into server hardware design, not added on later.

Dell Technologies has been a leader in using silicon-based security to secure firmware and protect sensitive user data in PowerEdge servers for the past two generations. The 14th, 15th, and 16th generations of PowerEdge product lines feature an enhanced Cyber Resilient Architecture that uses silicon-based Root-of-Trust to further harden server security including the following features:

- Cryptographically verified Trusted Booting that anchors end-to-end server and overall data center security. It includes features like silicon-based Root-of-Trust, digitally signed firmware, and automatic BIOS recovery.
- Secure Boot, which checks the cryptographic signatures of UEFI drivers and other code loaded before the operating system runs.
- iDRAC Credential Vault, a secure storage space for credentials, certificates, and other sensitive data that is encrypted with a silicon-based key that is unique for every server.
- Dynamic System Lockdown, a capability unique to PowerEdge servers, helps secure any system configuration and firmware from malicious or unintended changes while alerting users to any attempted system changes.
- Enterprise Key Management delivers a central key management solution to manage data-at-rest across the organization.
- System Erase, which allows users to easily retire or repurpose their 14th, 15th, and 16th generations of PowerEdge servers by securely and quickly wiping data from storage drives and other embedded nonvolatile memory.
- Supply Chain Security provides supply chain assurance by ensuring there is no product tampering or counterfeit components before shipping products to the customers.

The 14th, 15th, and 16th generations of PowerEdge servers, with their industry-leading security, form a trusted foundation for IT transformation on which customers can securely run their IT operations and workloads, and accelerate to Zero Trust adoption. Dell Technologies stops at nothing to help our customers build their breakthrough deployments. Our modern security approach ensures that an organization's environment is secure and resilient so that customers can focus on their core competencies, introduce their innovations, and advance human progress.



References

Dell Technologies documentation The following Dell Technologies documentation provides other information related to this document. Access to these documents depends on your login credentials. If you do not have access to a document, contact your Dell Technologies representative.

- iDRAC9 Security Configuration Guide
- Dell EMC Secured Component Verification Reference Guide for Servers
- <u>Understanding Confidential Computing with Trusted Execution Environments</u>
 <u>and Trusted Computing Base models</u>
- iDRAC9 System Lockdown: Preventing Unintended Server Changes
- Next Generation Dell PowerEdge Servers: Transition to Modern UEFI
- Dell EMC PowerEdge UEFI Secure Boot Customization: Reduce Attack Surface with Complete Control of Certificates
- Dell Technologies Supply Chain Security: Secured Component Verification for <u>PowerEdge</u>
- <u>Dell PowerEdge: iDRAC Automatic Certificate Enrollment</u>
- Improved iDRAC9 Security using TLS 1.3 over HTTPS on Dell PowerEdge Servers
- <u>A Partnership of Trust: Dell Supply Chain Security</u>
- <u>PowerEdge Advantages in your Zero Trust Journey Video</u>
- <u>AMD on PE Extending Data Protection to Data in Use Video</u>
- <u>AMD on PE Extended Boot Protection Video</u>
- <u>Zero Trust Architecture Video</u>
- Cyber Resilient Architecture Video
- <u>Secured Component Verification Video</u>
- <u>SEKM Video</u>
- IPv6 Direct from Development
- <u>iDRAC9 Virtual Power Cycle: Remotely power cycle Dell EMC PowerEdge</u>
 <u>Servers</u> -Direct from Development
- Transform Datacenter Analytics with iDRAC9 Telemetry Streaming
- <u>Configure iDRAC to use Active Directory Authentication (dell.com)</u>
- <u>Securing 14th Generation Dell EMC PowerEdge Servers with System Erase</u>
- (Direct from Development) Security in Server Design
- (Direct from Development) Cyber-Resiliency Starts At The Chipset And Bios
- <u>Factory Generated Default Password for iDRAC9 for Dell EMC 14th Generation</u>
 (14G) PowerEdge Servers

References

FIS. <u>331</u> Moy. 32

- <u>Dell EMC iDRAC Response to Common Vulnerabilities and Exposures (CVE)</u> <u>CVE-2017- 1000251 "BlueBorne"</u>
- <u>(Video) Secure Boot Configuration And Certificatemanagement Using</u> <u>RACADM- Video</u>
- Secure Boot Management on Dell EMC PowerEdge Servers
- <u>Signing UEFI images for Secure Boot feature in the 14th and 15th generation</u> and later Dell EMC PowerEdge servers
- <u>Rapid Operating System Recovery</u>
- <u>Managing iDRAC9 Event Alerts on 14th generation (14G) Dell EMC PowerEdge</u>
 <u>Servers</u>
- UEFI Secure Boot Customization
- iDRAC Overview
- OpenManage Console Overview
- OpenManage Mobile Overview
- <u>Motherboard Replacement</u>



D&LLTechnologies

Declaração do Fabricante

A **DELL COMPUTADORES DO BRASIL LTDA.** ("Dell"), inscrita no CNPJ sob o n. 72.381.189/0001-10, na qualidade de fabricante do(s) equipamento(s) de marca Dell (abaixo identificado(s)), ofertado(s) pela empresa Perfil Computacional, no certame licitatório n. PE 01/2025 promovido pelo PARANA PROJETOS, vem, através desta, declarar que:

- o(s) modelo(s) Dell PowerEdge R760xs possui(em) garantia de 60 meses, on-site, com atendimento telefônico 24 horas por dia, 7 dias na semana.

Declaramos, ainda, que:

- Os equipamentos por nós fabricados serão novos, sem uso e não são produtos descontinuados.

- A Perfil Computacional está autorizada a comercializar os equipamentos propostos para esse certame.

Eldorado do Sul, 28 de março de 2025



Dell Computadores do Brasil Ltda Luciano Valadares – Executivo de Contas

DELL Computadores do Brasil Ltda. Av. Industrial Belgraf, 400 . Eldorado do Sul / RS . Geral : 51 3481 5500 Fax : 51 3481 5458



D&LLTechnologies

Eldorado do Sul, 27 de março de 2025

À Perfil Computacional A/C Sr. Igor S. Reolon

Ref.: PARANÁ PROJETOS – Pregão Eletrônico 01/2025

DECLARAÇÃO TÉCNICA

DELL COMPUTADORES DO BRASIL LTDA. ("Dell"), inscrita no CNPJ/MF sob o n° 72.381.189/0001-10, com sede na Av. Industrial Belgraf, 400 – Medianeira – CEP 92990-000, Eldorado do Sul/RS, com o objetivo de complementar as informações que não constam no Catálogo Técnico Oficial do(s) produto(s) abaixo ofertado(s), vem, através da presente, declarar o que segue:

Objeto: Dell PowerEdge R760xs

- A BIOS é desenvolvida pela Dell Computadores do Brasil, detentora dos direitos copyright.

- Placa Mãe da fabricante DELL sendo projetada e desenvolvida especificamente para os modelos ofertado, não sendo uma placa de livre comercialização no mercado.

- Os manuais, drivers, firmwares e atualizações dos produtos Dell são disponibilizados no site do fabricante: www.dell.com.br/suporte.

- Todos os equipamentos e seus opcionais são integrados em fábrica, sem quaisquer adaptações.

- Os equipamentos ofertados são novos e estão em atual linha de produção.



Dell Computadores do Brasil Ltda Luciano Valadares – Executivo de Contas

DELL Computadores do Brasil Ltda. Av. Industrial Belgraf , 400 . Eldorado do Sul / RS . Geral : 51 3481 5500 Fax : 51 3481 5458



ProSupport Plus for Infrastructure

Introdução

A Dell Technologies¹ tem o prazer de fornecer o ProSupport Plus for Infrastructure (os "Serviços" ou "Serviços de suporte") em conformidade com esta Descrição de serviço ("Descrição de serviço"). A cotação, o formulário de pedido, outro formulário de fatura mutuamente acordado entre as partes ou a confirmação do pedido pela Dell Technologies (o "Formulário de pedido") incluirá o(s) nome(s) do(s) Produto(s)², o(s) Serviço(s) aplicável(is) e a(s) opção(ões) relacionada(s), caso existam. Para obter assistência adicional ou solicitar uma cópia do contrato vigente aplicável aos Serviços (o "Contrato"), entre em contato com um representante de vendas da Dell Technologies. Para obter uma cópia de seu contrato com o revendedor Dell Technologies aplicável, entre em contato com o revendedor.

O escopo deste Serviço

Os recursos deste Serviço incluem:

- Acesso por telefone 24 (vinte e quatro) horas por dia, 7 (sete) dias por semana (inclusive feriados)³ à central de especialistas globais da Dell Technologies, que é composta por engenheiros de nível sênior do ProSupport para assistência à solução de problemas de hardware e software.
- Envio no local de um técnico de campo especializado da Dell Technologies e/ou entrega de peças de substituição no Local de instalação ou em outra localização comercial indicada pelo Cliente e aprovada pela Dell Technologies, como detalhado no Contrato para resolver um problema no Produto.
- Acesso a um Service Account Manager (SAM).
- Substituição de peças, quando considerado necessário para resolver ou evitar um problema
- As plataformas AIOps do ProSupport incluem o APEX AIOps Infrastructure Availability, o TechDirect e o MyService 360, que são todos ativados pelo software de conectividade, como o Gateway de conexão segura (SCG), e oferecem benefícios não limitados a estes:
 - Detecção proativa de problemas e criação automatizada de casos
 - o Detecção preditiva de falhas de hardware
 - Criação de casos de autoatendimento
 - o Despacho de peças de autoatendimento
 - Avisos de segurança da Dell
 - Avaliação de segurança cibernética do APEX AIOps Infrastructure Availability

Consulte a tabela abaixo para obter mais informações.

Como entrar em contato com a Dell Technologies para solicitar um Serviço

Suporte on-line, por chat e por e-mail: O suporte da Dell Technologies por site, chat e e-mail está disponível para alguns produtos em <u>www.dell.com/contactus</u>.

Solicitação de suporte por telefone: Disponível 24x7 (inclusive em feriados). A disponibilidade pode variar fora dos Estados Unidos e limita-se a esforços comercialmente razoáveis, exceto quando especificado de outra forma neste documento. Acesse <u>www.dell.com/contactus</u> para obter uma lista de números de telefone aplicáveis à sua localização.

A tabela a seguir lista os recursos de serviço do ProSupport Plus for Infrastructure fornecidos de acordo com os termos da garantia e/ou manutenção da Dell Technologies. O ProSupport Plus for Infrastructure está disponível para dar suporte e manter:

1. Equipamentos da Dell Technologies identificados na <u>Tabela de garantia e manutenção de produtos da Dell Technologies</u> e/ou no Formulário de pedido como:

- Incluindo o ProSupport Plus for Infrastructure pela duração da garantia aplicável; ou

¹ "Dell Technologies", conforme usado neste documento, significa a entidade de vendas da Dell ("Dell") aplicável, especificada no Formulário de pedido da Dell, e a entidade de vendas da EMC ("EMC") aplicável, especificada no Formulário de pedido da EMC. O uso de "Dell Technologies" neste documento não indica alteração no nome legal da entidade Dell ou EMC com a qual você fez negócios.

² Conforme usado neste documento, "Produtos da Dell Technologies", "Produtos", "Equipamento" e "Software" significam o Equipamento e o Software da Dell Technologies identificados na <u>Tabela de garantia e manutenção de produtos da Dell Technologies</u> ou no seu Formulário de pedido, e "Produtos de terceiros" é definido no seu Contrato ou, na ausência de tal definição em seu Contrato, nos <u>Termos de venda comerciais da Dell Technologies</u> ou nos termos de venda locais da Dell Technologies, conforme aplicável. "Você" e "Cliente" referem-se à entidade nomeada como o comprador desses Serviços indicados no Contrato.

³ A disponibilidade varia de acordo com o país. Para obter mais informações, entre em contato com seu representante de vendas.



- Elegíveis ao upgrade para o ProSupport Plus for Infrastructure pela duração da garantia aplicável; ou
- Elegíveis ao ProSupport Plus for Infrastructure durante um período de manutenção subsequente.
- 2. Software da Dell Technologies identificado na <u>Tabela de garantia e manutenção de produtos da Dell Technologies</u> e/ou no Formulário de pedido como elegível ao ProSupport Plus for Infrastructure durante um período de manutenção.

RECURSO DO SERVIÇO	DESCRIÇÃO	PROSUPPORT PLUS — DETALHES DA COBERTURA
SUPORTE TÉCNICO GLOBAL	O Cliente entra em contato com a Dell Technologies por telefone ou pela interface Web 24x7 para informar um problema no Equipamento ou no Software. Os contatos por telefone serão encaminhados para um contato de suporte técnico remoto para resolução do problema. A criação automatizada de casos é disponibilizada quando a AlOps platform é configurada por meio do Gateway de conexão segura (SCG).	Incluso. • Para problemas de Severidade 1, os clientes obtêm gerenciamento de escala e procedimentos "CritSit" para situações críticas com a cobertura do Gerente de incidentes.
RESPOSTA NO LOCAL	A Dell Technologies envia profissionais autorizados ao Local de instalação para que trabalhem no problema após a Dell Technologies ter isolado o problema e confirmado a necessidade da Resposta no local.	 Incluída apenas para Equipamento. O objetivo inicial da Resposta no local é uma resposta de serviço em quatro horas depois que a Dell Technologies considerar necessário a Resposta no local. <u>Resposta no local</u> O técnico geralmente chega ao local em 4 horas após a conclusão da solução de problemas e o isolamento do problema. Disponível 24 (vinte e quatro) horas por dia, 7 (sete) dias por semana, inclusive feriados. Disponível nos locais definidos com resposta em 4 (quatro) horas. O estoque de componentes operacionais essenciais está a no máximo 4 horas de distância da localização do cliente, conforme determinado pela Dell Technologies. As peças não essenciais podem ser enviadas por meio de entrega noturna. A resposta no local não será aplicável ao software e poderá ser adquirida separadamente.
SUPORTE ESSENCIAL	Para problemas de Severidade 1, a Dell Technologies realiza a cobertura incluída, conforme considerado necessário pela Dell Technologies.	 Incluída apenas para Equipamento. Procedimentos de situação crítica — os problemas de Severidade 1 se qualificam à atuação rápida do Gerente de escala/resolução e à cobertura de incidentes "CritSit". Envio rápido: Despacho de engenheiro de campo sênior em conjunto com a solução de problemas por telefone. O engenheiro e a disponibilidade são determinados pela Dell conforme aplicável. Diagnóstico no local sob demanda da Dell quando a equipe do cliente não estiver disponível ou razoavelmente capaz de executar a solução de problemas no local. Somente para equipamento conectado ao gateway de conexão segura. A solicitação do cliente deve ser iniciada por meio de um chamado por telefone. Prioridade de produção em caso de situação crítica causada por um desastre natural. Em muitos casos, isso inclui a produção acelerada de um novo sistema da Dell Technologies.



RECURSO DO SERVIÇO	DESCRIÇÃO	PROSUPPORT PLUS — DETALHES DA COBERTURA
CHAMADA DE 6 HORAS PARA REPARO 24X7: RESPOSTA NO LOCAL EM 4 HORAS COM SERVIÇO DE REPARO DE HARDWARE DE 6 HORAS	O técnico deve chegar ao local em até 4 horas após o despacho e, em muitos casos, repara o hardware em até 6 horas após o despacho.	 Disponível 24 (vinte e quatro) horas por dia, 7 (sete) dias por semana, inclusive feriados. Para incidentes de Severidade 1, a Dell Technologies fará esforços comercialmente razoáveis para devolver o hardware à condição operacional em até 6 horas após o despacho Resposta em 4 horas e reparo em 6 horas após o despacho. Aplica-se somente a falhas ou reparos do Produto compatível coberto. O Suporte de software não está no escopo. Disponível para clientes a uma distância de até 80 quilômetros (50 milhas) do HUB de suporte designado da Dell Technologies. O Cliente deve ter uma versão compatível ativada e mantida do software de gateway de conexão segura.



RECURSO DO SERVIÇO	DESCRIÇÃO	PROSUPPORT PLUS — DETALHES DA COBERTURA
PLATAFORMAS AIOPS DO PROSUPPORT	AlOps é a inteligência artificial (IA) para as operações de TI. Ela se refere ao uso estratégico de tecnologias de IA, aprendizado de máquina (ML) e raciocínio de máquina (MR) que simplificam e facilitam processos e otimizam o uso dos recursos de TI do cliente.	Incluso. As plataformas AlOps do ProSupport incluem o APEX AlOps Infrastructure Availability, o TechDirect e o MyService 360, que são todos ativados pelo gateway de conexão segura da Dell Technologies, e oferecem benefícios não limitados a: • Detecção proativa de problemas e criação de casos • Detecção preditiva de falhas de hardware • Criação de casos de autoatendimento
		 Despacho de peças de autoatendimento Avisos de segurança da Dell Avaliação de segurança cibernética do APEX AlOps Infrastructure Availability
		O APEX AIOps Infrastructure Availability é um aplicativo de AIOps baseado em nuvem que fornece monitoramento e solução de problemas simples e proativos de sua infraestrutura de TI da Dell. Ele utiliza o aprendizado de máquina para monitorar e medir proativamente a integridade geral de servidores, armazenamento e dispositivos convergentes, hiperconvergentes, de proteção de dados e de rede por meio de análise preditiva, inteligente e abrangente. A análise preditiva da capacidade e do desempenho dos componentes, como unidades de estado sólido e memória, é habilitada por meio do uso do APEX AIOps Infrastructure Availability. O APEX AIOps Infrastructure Availability está disponível sem cobrança adicional para produtos com um contrato ProSupport (ou superior) válido. O APEX AIOps Infrastructure Availability é hospedado na nuvem privada da Dell Technologies, proporcionando a cada cliente um portal seguro e independente e garante que os clientes só poderão ver seu próprio ambiente.
		O MyService360 oferece visualizações e análises de dados de 360 graus em todo o data center, além de um histórico completo de serviços para produtos Dell Technologies. Os principais benefícios incluem:
		 Exibição proativa de incidentes e riscos críticos em tempo real Recomendações claras e prescritivas para simplificar a ação e o planejamento Tendências históricas e lógica analítica de dados Gerenciamento da integridade dos serviços para proteção de dados, armazenamento e sistemas convergentes
		O TechDirect permite o despacho automático de peças pelo cliente.
		O Gateway de conexão segura (SCG) é uma tecnologia de monitoramento empresarial fornecida como um equipamento e um aplicativo independente. Ele monitora seus dispositivos e detecta proativamente problemas de hardware que podem ocorrer. Dependendo do contrato de serviço, ele automatiza também a criação de solicitações de suporte para problemas detectados nos dispositivos monitorados.

Entre em contato com seu representante da Dell Technologies para saber mais sobre os produtos cobertos

necessário.

PEÇAS DE

REPOSIÇÃO



pelo cliente (CRUs). Consulte a Tabela de garantia e manutenção de produtos da Dell Technologies para obter uma lista das peças designadas como CRUs para Equipamento específico ou entre em contato com a Dell Technologies para obter mais detalhes.

Se a Dell Technologies instalar a peça de substituição, ela providenciará a devolução da peça substituída a uma instalação da Dell Technologies. Se o Cliente instalar a CRU, ele será o responsável pela devolução da CRU substituída à instalação designada pela Dell Technologies. Se um cliente precisar de assistência para substituir uma CRU, a Dell Technologies poderá ajudar remotamente e/ou enviar um técnico no local para ajudar na substituição.

Caso o técnico da Dell conclua, durante o diagnóstico, que o reparo pode ser realizado com uma peça designada como CRU ou se o Cliente optar por despachar automaticamente uma peça designada como CRU, a Dell enviará essa peça diretamente ao Cliente.

Se o técnico da Dell determinar que o Produto com Suporte deve ser substituído por inteiro, a Dell Technologies reserva-se o direito de enviar ao Cliente uma unidade de substituição inteira. As substituições de unidades inteiras podem não estar disponíveis em estoque para tempos de resposta no mesmo dia, e pode haver prazos estendidos



		para a chegada de uma substituição de unidade inteira em sua localização, dependendo de onde você está localizado e do tipo
		de Produto que está sendo substituído.
SUBSTITUIÇÃO	Se, antes de atingir seu Nível de	Incluída para Produtos de armazenamento e infraestrutura
PROATIVA DE	Resistência, uma unidade de estado sólido atingir o Limite de Nível de	hiperconvergente/convergente.
UNIDADE SSD	Resistência (conforme determinado	O objetivo de resposta baseia-se nos recursos de serviços aplicáveis
	pela Dell) ou superior, o Cliente será	de Entrega de peças de substituição e Resposta no local, detalhados
	elegivel para receber uma unidade	acima. O cliente deve ativar e manter as versoes atualmente
	de estado solido de substituição. "Nível de resistêncie" significa a vida	TI (implementado como um gatoway do conovão cogura) duranto
	útil média de uma SSD elegível	o período de suporte aplicável. A ativação do software de
	"Limite de nível de resistência"	conectividade, conforme aplicável, é um pré-requisito para
	é o ponto na vida útil da SSD	esses outros recursos de serviço de renovação.
	(conforme determinado pela Dell)	
	no qual a unidade se torna elegível	Unidades pré-criadas não são elegíveis para Substituição proativa
	para substituição, por exemplo, ao	da unidade SSD pela Dell Technologies.
	alingir 95% do nivel de resistência.	
	variarão.	
DIREITOS	A Dell Technologies oferece os	Incluso.
A NOVAS	offware conforme as disponibiliza	
VERSÕES	software comorne as disponibiliza.	
DE SOFTWARE		
		Software de embiente energeienel (OE) de Equinemente
INSTALAÇÃO DE	a instalação remota das	Software do ambiente operacional (OE) do Equipamento
NOVAS VERSOES	novas Versões de software.	Incluído apenas no equipamento de armazenamento quando o
DE SOFTWARE		software associado do ambiente operacional incorporado está coberto
		por uma garantia da Dell ou por um contrato de manutenção em vigor
		da Dell. O software do ambiente operacional do equipamento é
		definido como programação e/ou microcódigo do software de interface
		do usuario necessario para permitir a administração, o controle e a
		equipamento não pode operar
		O cliente tem direito à instalação remota das atualizações de
		software do de com uma versão compativer ativada e mantida do
		Consulte a <u>Tabela de garantia e manutenção de produtos</u> para
		produtos elegíveis.
		Outros softwares (não OE)
		O Cliente realiza a instalação de novas Versões de Software, a menos
		que seja considerada necessaria pela Dell Technologies.
	Certos Produtos acionam um contato	Incluso para Produtos que têm SCG ou outras ferramentas de
DEMOTO	com a Dell Technologies de modo	conectividade habilitadas pela Dell, para tecnologia de monitoramento
	automático e independente para	remoto disponível na Dell Technologies.
E REPARO 24X7	fornecer dados que auxiliem	Voia os detalhos sobre o forremento SCC poimo na osoão Blataforma
	o problema.	
	- F. Solina.	
	A Dell Technologies acessa os	Quando a Dell Technologies for notificada sobre um problema, os
	Produtos remotamente, caso	mesmos objetivos de resposta do Suporte técnico global e da
	adicionais e dar suporte remoto.	Resposta no local serao aplicavels, conforme descrito anteriormente.
	•	

Dell Technologies ProSupport Plus for Infrastructure | v11 | May 2024

Inserido ao protocolo **22.951.206-4** por: **Pedro Henrique Golin Linhares** em: 31/03/2025 12:49. A autenticidade deste documento pode ser validada no endereço: https://www.eprotocolo.pr.gov.br/spiweb/validarDocumento com o código: **692a5ef82beb173ed33dad4ee6505546**.



SERVICE ACCOUNT MANAGER ("SAM") O SAM atribuído do ProSupport Plus for Infrastructure é um recurso remoto que oferece uma ampla variedade de recursos e funcionalidades de gerenciamento de sistema, ambiente e conta, criados para reduzir o tempo de inatividade e melhorar a experiência geral de suporte da Dell Technologies.

Estão incluídos no Serviço: Assistência de integração:

- Verificando a exatidão das informações relevantes de suporte ao cliente, como nome da conta, endereço etc.
- Fornecendo transferências de conhecimento, tais como entrar em contato com a Dell Technologies para abrir chamados e o uso de ferramentas e tecnologias de suporte da Dell Technologies
- Designando o cronograma para entregas do SAM, como relatórios e análises de serviço

<u>Relatórios mensais</u>: Geração de relatórios e recomendações sobre sistemas qualificados, incluindo:

- Resumo das chamados abertos e fechados por mês;
- Verificação de versões de software do sistema atualmente instaladas com base em recomendações de código de destino; e
- Status do contrato, incluindo datas de início/fim e outros detalhes básicos do contrato.

Para permitir totalmente a geração de relatórios mensais, as tecnologias de conectividade da Dell Technologies, como o gateway de conexão segura, devem ser instaladas com as opções apropriadas de coleta de log ativadas.

Análise de serviço: O SAM fornece uma análise de serviço dos detalhes no relatório de serviço. Prazo, cronograma e outros temas a serem analisados serão determinados entre o SAM e o Cliente durante a integração.

Manutenção do sistema: Para ativos qualificados, o SAM auxiliará o Cliente na coordenação da entrega de eventos de manutenção do sistema dentro da janela de manutenção do Cliente. Veja abaixo informações adicionais.

<u>Suporte de escalonamento da</u> <u>Dell Technologies</u>: atuar como elo de Serviços para coordenar os recursos necessários e resolver questões individuais de Severidade 1 ou problemas mais sistêmicos. Incluído nos Produtos cobertos pelo serviço ProSupport Plus for Infrastructure ou no contrato de manutenção em vigor durante o horário comercial local normal da Dell Technologies, que pode variar de acordo com a região e o país, exceto em feriados locais e da Dell Technologies. Consulte os detalhes adicionais da cobertura abaixo.

A Dell Technologies é responsável por realizar apenas as atividades e tarefas do SAM expressamente especificadas neste documento. Todas as outras tarefas, atividades e serviços estão fora do escopo.



DEFINIÇÕES DE NÍVEL DE SEVERIDADE

SEVERIDADE 1 Crítica — incapaz de executar funções essenciais aos negócios e requer resposta imediata.

SEVERIDADE 2: Alta — capaz de executar funções de negócios, mas o desempenho/capacidades estão degradados ou gravemente limitados.

SEVERIDADE 3 Média/baixa - mínimo ou nenhum impacto nos negócios.

Informações adicionais sobre o Suporte essencial

A Dell se reserva o direito de recusar o serviço se a Dell Technologies acreditar que o Cliente está usando incorretamente ou em demasia o recurso On-Demand Onsite Diagnosis para problemas críticos (por exemplo, quando o cliente tiver uma equipe disponível para solucionar o problema no local ou as solicitações do Cliente para as visitas de diagnóstico no local excederem significativamente as taxas de falha padrão, em decorrência dos exemplos a seguir, dos componentes e sistemas envolvidos). Se a Dell Technologies determinar (a critério exclusivo da Dell Technologies) que um Cliente está utilizando abusivamente o Serviço, a Dell Technologies se reservará o direito de negar o serviço.

Exclusões

As seguintes atividades não estão incluídas no escopo desta Descrição de serviço:

- Desinstalação, reinstalação ou configuração de produtos, software ou aplicativos
- Remoção de Produto desinstalado do local do Cliente
- O software de servidor/armazenamento/rede não é classificado como equipamento
- Solução de problemas de software do ambiente operacional que não reabilitarão o produto a um estado funcional (por exemplo, a consultoria, o ajuste de desempenho, a configuração, a criação de scripts ou a análise comparativa não está incluído)
- Serviços necessários devido à falta de manutenção do software e dos Produtos compatíveis com o suporte em qualquer nível mínimo de versão especificado, conforme estabelecido no <u>Documento de código de referência</u>.
- Serviços necessários devido à não execução de qualquer correção, reparo, patch ou modificação do sistema fornecidos pela Dell Technologies ou à falha do Cliente em tomar medidas de contenção previamente recomendadas pela Dell Technologies, como avisos de segurança ou atualizações de correção essenciais comunicados e que o cliente não implementou
- Serviços que, segundo a Dell Technologies, são necessários devido ao tratamento ou uso inadequado dos produtos ou equipamentos
- Personalização do servidor ou do dispositivo de armazenamento do Cliente, exceto quando expressamente declarado nesta
 Descrição de serviço
- Qualquer recuperação ou transferência de dados ou aplicativos
- Serviço de garantia ou suporte para sistemas, software ou componentes adicionais que não sejam da Dell Technologies
- Serviços que, segundo a Dell Technologies, são necessários devido a tentativas não autorizadas por pessoal terceirizado de instalar, reparar, manter ou modificar hardware, firmware ou software
- Instalação da impressora de rede ou mapeamento do compartilhamento de arquivos de rede
- Configuração de qualquer tipo para servidor, armazenamento, rede ou roteador
- Serviços de rede, incluindo conexão de um sistema a uma rede (diferente de uma LAN Ethernet)
- Qualquer atividade não estabelecida especificamente nesta Descrição de serviço.

Esta Descrição de serviço não concede ao Cliente nenhuma garantia adicional às garantias fornecidas segundo os termos do contrato principal de serviços ou do Contrato, conforme aplicável.



Responsabilidades do Cliente pelo recurso do serviço do SAM

O fornecimento por parte da Dell Technologies do recurso de serviço SAM detalhado acima dependerá do cumprimento do Cliente das seguintes responsabilidades:

- Disponibilizar as janelas de manutenção do sistema apropriadas para o SAM conforme considerado necessário pela Dell Technologies.
- Garantia de que todos os requisitos ambientais, técnicos e operacionais sejam atendidos.
- Fornecer ao SAM acesso oportuno a (a) pelo menos um contato técnico com responsabilidades de administração do sistema e privilégios apropriados de acesso ao sistema/informações, e (b) especialistas no assunto, sistemas e redes aplicáveis (incluindo, entre outros, acesso a sistemas remotos/rede) conforme considerado necessário pela Dell Technologies.
- Assumir toda a responsabilidade pela conectividade de rede, pelo desempenho e pelos problemas de configuração.
- Verificar se as localizações dos Equipamentos estão preparadas antes do início do ProSupport Plus for Infrastructure.
- Garantir que o produto esteja conectado ao gateway de conexão segura com as opções apropriadas de coleta de log ativadas.

Outras informações importantes sobre o recurso de serviço do SAM

- O serviço do SAM está disponível durante o horário comercial normal. O horário comercial é definido pelo local onde o SAM reside e pode variar de acordo com a região e o país. A critério da Dell Technologies e quando ela considerar necessário, os serviços do SAM podem ser realizados no local.
- O suporte depois do expediente pode ser fornecido por outros recursos da Dell Technologies a critério da Dell Technologies.
- O local do SAM será atribuído por acordo mútuo durante a integração, com base na área de serviço preferencial do cliente e na disponibilidade da equipe da Dell.

MANUTENÇÃO PROATIVA DO SISTEMA PROSUPPORT PLUS FOR INFRASTRUCTURE

A Manutenção do sistema do ProSupport Plus fornece aos clientes da Dell Technologies eventos de manutenção remota proativa e planejada que ocorrem durante a vigência do contrato de serviço em dispositivos cobertos pelo ProSupport Plus for Infrastructure e monitorados com a tecnologia de conectividade aplicável da Dell Technologies, como o gateway de conexão segura, conforme aplicável. Esses eventos de manutenção do sistema ajudam a manter o desempenho e podem reduzir a probabilidade de incidentes futuros devidos a versões incompatíveis de hardware, software, BIOS e firmware. Os eventos de manutenção de sistema proativa e planejada são coordenados entre os clientes, o SAM e o pessoal de suporte da Dell Technologies. A entrega desses eventos geralmente está disponível 24x7x365, mas pode estar sujeita à mútua disponibilidade de recursos do cliente e da Dell Technologies. A Dell Technologies recomenda que a manutenção de sistema proativa e planejada ocorra duas vezes por ano. Certos produtos podem ter limitações no número de vezes que a manutenção de sistema proativa e planejada pode ser realizada por ano. Consulte o representante de vendas ou o SAM designado para obter uma lista de Produtos com Suporte e saber guais são as limitações aplicáveis.

Não incluídos na manutenção do sistema do ProSupport Plus for Infrastructure

- Atualizações em dispositivos interconectados não cobertos por um contrato de suporte atual do ProSupport Plus for Infrastructure.
- Atualizações em qualquer software sem o direito correspondente a tais atualizações de acordo com um contrato de suporte de software apropriado por parte da Dell Technologies ou de terceiros para alguns Produtos de Terceiros.
- Upgrades de sistema operacional e criação de patches de hypervisor ou outro suporte de engenharia ou de desenvolvimento de software relacionado.
- Criação de patches de aplicativo.
- Prestação de serviços de manutenção no local.
- Desinstalação ou instalação de hardware adicional, ou tarefas de configuração.
- Instalação ou configuração de software não especificamente listado nesta descrição de serviço.
- Ajuste de desempenho do aplicativo.
- Remoção ou identificação de vírus, spyware, ou malware.
- Todas as outras atualizações ou outras atividades não especificamente documentadas dentro desta Descrição do serviço.


Informações adicionais importantes sobre a manutenção do sistema do ProSupport Plus for Infrastructure

- Durante o evento de manutenção, upgrades podem causar uma perda temporária de conectividade com outros dispositivos conectados.
- Após a conclusão da atualização, os dispositivos conectados talvez precisem ser reiniciados e a conectividade verificada.
- Os sistemas a serem submetidos a upgrade devem ser disponibilizados para a Dell Technologies ou para os agentes autorizados pela Dell Technologies durante a janela de manutenção acordada.
- Dependendo dos sistemas a serem atualizados, pode ser necessário disponibilizar um sistema ou recurso de gerenciamento de sistema adicional.
- Dependendo dos sistemas a serem submetidos a upgrade, pode ser necessário fornecer direitos de administração apropriados para o dispositivo à Dell Technologies ou aos agentes autorizados da Dell Technologies.
- O cliente é responsável por ter e manter todos os requisitos de licença relativos às atualizações de hardware e software.
- Caso essa atualização do Software em um Produto qualificado possa reduzir ou afetar o desempenho de outro Produto não qualificado, a Dell Technologies, em consulta ao cliente, pode optar por não dar andamento à atividade de manutenção do sistema até que a situação seja resolvida.
- Para permitir totalmente a Manutenção de Sistema do ProSupport Plus for Infrastructure, a tecnologia de conectividade aplicável, como o gateway de conexão segura, deve ser instalada com as opções de coleta de log ativadas.
- A instalação de novas Versões de software para os sistemas de armazenamento de avançados e os sistemas de infraestrutura convergente e hiperconvergente da Dell Technologies determinadas durante a manutenção do sistema, inclusive, entre outras, versões de software publicadas nas matrizes de configuração de interoperabilidade aplicáveis (a Matriz de suporte simples ou a Matriz de certificação de versão da Dell Technologies), pode exigir a compra de um projeto de serviços separado da Dell Technologies para obter mais informações.

Entrega de relatórios a revendedores autorizados da Dell Technologies para Clientes usuários finais que compram de revendedores autorizados da Dell Technologies

O relatório do SAM do ProSupport Plus será fornecido pela Dell Technologies apenas ao **revendedor ou distribuidor autorizado da Dell Technologies (conforme aplicável) identificado na cotação da Dell Technologies (coletivamente, o "Revendedor"). O Revendedor reconhece e concorda que obteve consentimento apropriado dos clientes usuários finais do Revendedor ("Usuários Finais") para receber os Relatórios do SAM do cliente Usuário Final**. A não ser que tenha sido notificado em contrário pelo Revendedor, a Dell Technologies distribuirá os relatórios do SAM do ProSupport Plus para o endereço e as informações de contato do Revendedor fornecidos à Dell Technologies quando o Serviço foi adquirido ou para as informações de contato contidas nos registros atuais de vendas e serviços da Dell Technologies como determinado por ela. Os relatórios do SAM fornecidos ao Revendedor não serão categorizados por/para Usuários Finais específicos. Opções de relatórios personalizados podem estar disponíveis para compra separada por um custo adicional.

ASSISTÊNCIA COLABORATIVA

Se o Cliente abrir um chamado e a Dell Technologies determinar que o problema surge com um produto de fornecedor terceirizado elegível comumente utilizado em conjunto com os Produtos cobertos por uma garantia ou por um contrato de manutenção atual da Dell Technologies, a Dell Technologies fará esforços para fornecer Assistência colaborativa, em que a Dell Technologies: (i) serve como ponto único de contato até que os problemas sejam isolados; (ii) entra em contato com o fornecedor terceirizado; (iii) fornece documentação do problema; e (iv) continua a monitorar o problema e obtém o status e planos de resolução do fornecedor (quando razoavelmente possível).

Para ter direito à Assistência colaborativa, o Cliente deve ter os contratos de suporte ativos e os direitos apropriados diretamente com o respectivo fornecedor terceirizado e a Dell Technologies ou com um revendedor autorizado Dell Technologies. Depois que o problema é isolado e informado, o fornecedor terceirizado torna-se o único responsável por oferecer todo o suporte, seja ele técnico ou não, em conexão com a resolução do problema do Cliente. A Dell Technologies NÃO SE RESPONSABILIZA PELO DESEMPENHO DE PRODUTOS OU SERVIÇOS DE OUTROS FORNECEDORES. Uma lista de parceiros de Assistência colaborativa pode ser encontrada na Lista de assistência colaborativa. Observe que os produtos cobertos pelo suporte de terceiros podem ser alterados a qualquer momento, sem qualquer aviso prévio aos Clientes.



SUPORTE DE SOFTWARE DO SISTEMA DA DELL TECHNOLOGIES

O suporte de software da Dell Technologies incluído no ProSupport Plus for Infrastructure destina-se a alguns Produtos de terceiros, incluindo determinados aplicativos de usuário final, sistemas operacionais, hypervisores e firmware, quando esses Produtos de terceiros 1) estão instalados e funcionando e são usados com os Produtos no momento em que o suporte é solicitado, 2) estão cobertos por um termo de servico de suporte e manutenção do ProSupport Plus for Infrastructure e 3) têm os direitos e contratos de suporte ativos apropriados diretamente com o respectivo editor de Produtos de terceiros. Esse nível de suporte é fornecido no Equipamento qualificado do ProSupport Plus for Infrastructure, independentemente de como o software qualificado foi adquirido e licenciado, mas o Cliente é responsável por garantir que esse software qualificado tenha sido adquirido e licenciado corretamente de acordo com o editor. O Cliente é o único responsável por corrigir quaisquer problemas com licenças e compras de software qualificado para poder receber esses Serviços a qualquer momento durante o período de cobertura. Os aplicativos de software elegíveis podem ser encontrados na Lista abrangente de suporte de software. Os produtos suportados de terceiros podem ser alterados a qualquer momento, sem aviso aos Clientes. As situações que dão origem às perguntas do cliente devem ser reprodutíveis em um sistema único, que pode ser físico ou virtual. O cliente compreende e aceita que as soluções de alguns problemas que dão origem ao chamado do Cliente podem não ser fornecidas pelo editor do software em questão (incluindo, entre outras, situações em que o editor não está mais fornecendo suporte ou manutenção ao software em questão por qualquer motivo) ou podem exigir o suporte adicional do editor, inclusive a instalação de software adicional ou outras alterações nos produtos. O cliente aceita que, nessas situações nas quais o editor do software em questão não fornece uma resolução, a obrigação da Dell Technologies de fornecer suporte ao cliente será plenamente cumprida.

Termos e condições adicionais aplicáveis aos usuários finais que compram Produtos de um OEM

O "OEM" é um revendedor que vende os Produtos compatíveis como um fabricante de equipamento original que está adquirindo os Produtos e Serviços da Dell Technologies do grupo de negócios OEM Solutions (ou seu sucessor) para um projeto de OEM. O OEM normalmente incorpora ou agrupa esses Produtos da Dell Technologies com hardware, software ou outra propriedade intelectual pertencente ao Cliente OEM, resultando em um sistema ou solução especializado com a funcionalidade específica da tarefa ou do setor (esse sistema ou solução sendo uma "Solução de OEM") e revende essa Solução de OEM com sua própria marca. Com relação a OEMs, o termo "Produtos com Suporte" inclui os Produtos com Suporte da Dell Technologies que são fornecidos sem a marca Dell Technologies (ou seja, sistema pronto de OEM sem marca), e "Usuário final" significa você ou qualquer entidade que adquire uma Solução de OEM para sua própria utilização final e não para revenda, distribuição ou sublicenciamento a terceiros. É responsabilidade do OEM oferecer ao Usuário Final solução de problemas de primeiro nível. Um diagnóstico inicial apropriado de melhor esforço deverá ser realizado pelo OEM antes de a chamada prosseguir para a Dell Technologies. Esse OEM permanece com a responsabilidade de fornecer a solução de problemas inicial, mesmo quando o Usuário final contata a Dell Technologies para solicitar o serviço. Além disso, se um Usuário final contatar a Dell Technologies para solicitar o serviço. Além disso, se um Usuário final contatar seu OEM, a Dell Technologies pedirá que o Usuário final contate o OEM para receber uma solução de problemas de primeiro nível antes de entrar em contato com a Dell Technologies.

Dell Technologies ProSupport Plus for Infrastructure sobre peças não padrão em produtos personalizados de servidor

Os reparos e substituições de peças não padrão ou únicas ("Serviços de Suporte de Componente Não Padrão") são um serviço de substituição de valor agregado que complementa a garantia do Produto PowerEdge do Cliente em componentes padrão da Dell Technologies em uma configuração padrão e que exigem substituição devido a defeitos de fabricação ou de material ("Reparos na Garantia"). Firmware/software com a marca Dell Technologies NÃO estão disponíveis para "Componentes Não Padrão", e o Cliente deve usar utilitários fornecidos pelo fabricante para monitorar e/ou atualizar o componente. O Cliente trabalhará diretamente com o fabricante para resolver qualquer problema de qualidade relativo a software/firmware, utilitários e hardware. A Dell Technologies prestará Serviços de Suporte Não Padrão para substituir peças não padrão ou únicas que o Cliente preveja e garanta que estarão disponíveis como definido acima e desde que o Cliente tenha feito os arranjos correspondentes para ajudar a Dell Technologies no processo de pedidos para estoque de serviço a fim de facilitar a atividade de reparo. Desde que o Cliente tenha uma previsão exata das necessidades de estoque, a Dell Technologies trocará a peca que apresente defeito de acordo com o tempo de resposta aplicável do Cliente para Reparos na Garantia e instalará a peça de substituição no Produto do Cliente, mas o Cliente reconhece e concorda que a Dell Technologies não é responsável perante ele por garantir a disponibilidade das peças. As peças e a resposta em campo no mesmo dia (por exemplo, 4 horas) podem não estar disponíveis para substituição de componentes "não padrão". Nesses casos, a Dell Technologies definirá como padrão o serviço no próximo dia útil (ou seja, após o diagnóstico e a solução do problema por telefone, uma peça ou um técnico geralmente será despachado no próximo dia útil). As peças de substituição podem ser novas ou reformuladas conforme permitido pelas leis locais, e a realização de reparos e substituições de Serviços de Suporte de Componente Não Padrão pode exigir que a Dell Technologies utilize garantia e/ou serviços de manutenção de um fabricante/editor terceirizado, e o Cliente concorda em prestar assistência à Dell Technologies e em fornecer todo o material solicitado por qualquer fabricante ou editor terceirizado para facilitar a utilização da garantia e/ou dos serviços de manutenção do terceiro correspondente.



Os testes de engenharia da configuração resultante realizados pela Dell Technologies e baseados em uma declaração de trabalho SOW separada, por exemplo, testes executados depois da instalação das peças não padrão ou exclusivas para uma configuração que utiliza o software solicitado pelo cliente, são uma atividade pontual a ser realizada uma vez, e não continuamente, e os serviços de suporte de componente não padrão estão disponíveis apenas para a configuração específica conforme definida pelo cliente e testada pela Dell Technologies. A Dell Technologies comunicará a configuração de hardware exata testada, inclusive os níveis de firmware. Depois que os testes de engenharia forem concluídos, a Dell Technologies fornecerá os resultados por meio de relatórios com a indicação de Aprovado/Reprovado. A Dell Technologies envidará esforços comercialmente razoáveis para dar suporte ao reconhecimento e à operação do componente não padrão no Produto da Dell Technologies, no entanto, a modificação de utilitários padrão da Dell Technologies (inclusive BIOS, IDRAC e software de conectividade) não será aceita. O Cliente será responsável por trabalhar diretamente com o fabricante para resolver qualquer problema do componente não padrão que surja durante o teste de engenharia (inclusive problemas de qualidade, software, firmware ou especificações/limitações de hardware). Os testes de engenharia adicionais da Dell Technologies depois que o Cliente tiver recebido um relatório com a indicação de APROVADO exigirão uma nova SOW e tarifas de engenharia não recorrentes associadas, inclusive em qualquer teste de engenharia solicitado em conexão com

Responsabilidades Gerais do Cliente

Autoridade para Conceder Acesso. O Cliente declara e garante que obteve permissão tanto para o Cliente quanto para a Dell Technologies de acessar e usar, seja remota ou pessoalmente, o software, o hardware, os sistemas licenciados ou pertencentes ao Cliente e os dados neles contidos, além de todos os componentes de hardware e software neles incluídos, para a finalidade de prestar esses Serviços. Se o Cliente ainda não tiver essa permissão, caberá a ele obtê-la, às suas custas, antes de solicitar que a Dell Technologies preste esses serviços.

Não solicitação. Quando permitido pela lei, o cliente não solicitará – sem a aprovação prévia por escrito da Dell Technologies, por um período de dois anos a partir da data indicada no seu Formulário de pedido – direta ou indiretamente a contratação de qualquer funcionário da Dell Technologies com quem ele tenha entrado em contato devido à prestação do Serviço pela Dell Technologies, desde que, no entanto, anúncios gerais e outras formas amplas de contratação não constituam solicitação direta ou indireta nos termos deste instrumento e você tenha permissão para propor a contratação de qualquer funcionário que tenha sido demitido ou tenha pedido demissão do emprego na Dell Technologies antes do início das conversas sobre emprego com você.

Cooperação do Cliente. O Cliente compreende que, sem cooperação rápida e adequada, a Dell Technologies não conseguirá realizar o Serviço ou, se realizado, o Serviço poderá ser significativamente alterado ou atrasado. Sendo assim, de maneira imediata e aceitável, o Cliente cooperará integralmente com a Dell Technologies, conforme necessário, para que ela possa realizar o Serviço. Se o Cliente não cooperar de forma razoavelmente adequada, de acordo com o disposto acima, a Dell Technologies não será responsável por nenhuma falha na realização do Serviço, e o Cliente não terá direito a reembolso.

Obrigações no Local. Quando os Serviços exigirem trabalho no local, o Cliente fornecerá (sem custo para a Dell Technologies) acesso livre, seguro e suficiente às instalações e ao ambiente do Cliente, incluindo amplo espaço de trabalho, eletricidade, equipamento de segurança (se aplicável) e linha telefônica local. O Cliente deverá fornecer também (sem custo para a Dell Technologies) um monitor ou uma tela, um mouse (ou dispositivo de apontamento eletrônico) e um teclado, caso o sistema já não inclua esses itens.

Backup dos Dados. O Cliente fará um backup de todos os dados, software e programas existentes em todos os sistemas afetados antes e durante a prestação deste Serviço. O Cliente deve fazer cópias de backup regulares dos dados armazenados em todos os sistemas afetados, como precaução contra possíveis falhas, alterações ou perdas de dados. A Dell Technologies não será responsável pela restauração ou reinstalação de nenhum programa nem de dados. Exceto quando proibido pelas leis locais aplicáveis, a Dell Technologies não terá responsabilidade por perda de dados em relação a:

- 1. Quaisquer informações confidenciais, de propriedade exclusiva ou pessoais;
- 2. Dados, programas ou software perdidos ou corrompidos;
- 3. Mídia removível perdida ou corrompida;
- 4. Perda do uso de um sistema ou rede; e/ou
- 5. Qualquer ato ou omissão, inclusive negligência, da Dell Technologies ou de um prestador de serviços terceirizado.

Garantias de Terceiros. Estes serviços podem exigir que a Dell Technologies tenha acesso a hardware ou software que não tenha sido produzido ou vendido por ela. As garantias de alguns fabricantes podem ser anuladas se a Dell Technologies ou outra pessoa que não o fabricante trabalhar no hardware ou no software. O Cliente garantirá que a execução dos serviços pela Dell Technologies não afetará tais garantias ou, se isso ocorrer, que o resultado será aceitável para o Cliente. A Dell Technologies não se responsabiliza por garantias de terceiros nem por efeitos que os serviços possam ter nessas garantias.



Manter as versões do Software e Serviço. O Cliente deverá manter o software e os Produtos compatíveis com o suporte, de acordo com as os níveis de versão mínimos especificados pela Dell Technologies no Documento de código de referência. O Cliente deverá também garantir a instalação dos níveis de versão mínimos de software e firmware nas peças de reposição, patches, atualizações de software ou versões subsequentes, conforme indicado pela Dell, para manter os direitos dos Produtos compatíveis com o suporte à realização do Serviço. A Dell Technologies reserva-se o direito, a seu critério exclusivo, de negar suporte a quaisquer software e Produtos compatíveis com o suporte que não atendam aos níveis mínimos de versão especificados pela Dell Technologies, conforme especificado no Documento de código de referência.

Termos e Condições dos Serviços

Esta Descrição de Serviço é acordada entre você, o Cliente ("você" ou "Cliente") e a Dell Technologies. Este Serviço está sujeito e é regido pelo Acordo do Cliente com a Dell Technologies.

Os produtos ou serviços obtidos de qualquer revendedor Dell Technologies serão regidos exclusivamente pelo contrato firmado entre o comprador e o revendedor. Esse contrato pode fornecer termos iguais aos deste documento ou dos termos on-line abaixo. O revendedor pode firmar acordos com a Dell Technologies para realizar serviços de garantia e/ou manutenção para o comprador em nome do revendedor. Os clientes e os revendedores que realizam serviços de garantia e/ou manutenção ou serviços profissionais devem ser adequadamente treinados e certificados. A realização de qualquer serviço por Clientes, revendedores ou terceiros não treinados/não certificados pode resultar em tarifas adicionais se o suporte da Dell Technologies for necessário em resposta à prestação de serviços por terceiros. Entre em contato com o revendedor ou o representante de vendas local da Dell Technologies para obter mais informações sobre a prestação de serviços de garantia e manutenção da Dell Technologies para Produtos adquiridos de um revendedor.

Na ausência de tal acordo autorizando explicitamente este Serviço e dependendo da localização do Cliente, este Serviço estará sujeito e será regido pelos Termos comerciais de venda da Dell ou pelo contrato de revenda ao qual a tabela a seguir fizer referência. Consulte a tabela abaixo, que mostra o URL aplicável à localização do Cliente onde seu contrato está disponível. As partes confirmam que leram e concordam com o cumprimento destes termos on-line.

Lecelizeeãe	- Termos e condiç	ões aplicáveis à compra dos Serviços
- Localização do Cliente	 Clientes que compram os Serviços diretamente 	 Clientes que compram os Serviços por meio de um revendedor Authorized
- Estados Unidos	- <u>Dell.com/CTS</u>	- <u>Dell.com/CTS</u>
- Canadá	- <u>Dell.ca/terms</u> (inglês) <u>Dell.ca/conditions</u> (francês do Canadá)	- <u>Dell.ca/terms</u> (inglês) <u>Dell.ca/conditions</u> (francês do Canadá)
- Países da América Latina e do Caribe	Site <u>Dell.com</u> local específico do país ou <u>Dell.com/servicedescriptions</u> .*	As Descrições de serviço e outros documentos de serviço da Dell Technologies que você possa receber do vendedor não constituirão um contrato entre você e a Dell Technologies, servindo apenas para descrever o conteúdo do Serviço que você está adquirindo desse vendedor, suas obrigações como destinatário do Serviço e os escopos e limitações de tal Serviço. Como consequência e neste contexto, qualquer referência a "Cliente" nesta Descrição de Serviço e em outros documentos de serviço da Dell Technologies será entendida como uma referência a você, enquanto as referências à Dell Technologies devem ser entendidas como uma referência à Dell Technologies como um provedor de serviços que está fornecendo o Serviço em nome do seu vendedor. Você não terá um relacionamento contratual direto com a Dell Technologies em relação ao Serviço descrito neste documento. Para evitar dúvidas, qualquer condição de pagamento ou contratual que seja, por natureza, diretamente relevante apenas entre um comprador e um vendedor não será aplicável a seu caso e seguirá o acordado entre você e seu vendedor.
- Ásia- Pacífico e Japão	Site <u>Dell.com</u> local específico do país ou <u>Dell.com/servicedescriptions</u> .*	As Descrições de Serviço e outros documentos de serviço da Dell Technologies que você possa receber do vendedor não constituem um contrato entre você e a Dell Technologies, servindo apenas para descrever o conteúdo do Serviço que você está adquirindo desse vendedor, suas obrigações como destinatário do Serviço e os escopos e limitações de tal Serviço. Como consequência e neste contexto, gualguer referência a "Cliente"



		nesta Descrição de Serviço e em outros documentos de serviço da Dell Technologies será entendida como uma referência a você, enquanto as referências à Dell Technologies devem ser entendidas como uma referência à Dell Technologies como um provedor de serviços que está fornecendo o Serviço em nome do seu vendedor. Você não terá um relacionamento contratual direto com a Dell Technologies em relação ao Serviço descrito neste documento. Para evitar dúvidas, qualquer condição de pagamento ou contratual que seja, por natureza, diretamente relevante apenas entre um comprador e um vendedor não será aplicável a seu caso e seguirá o acordado entre você e seu vendedor.
- Europa, Oriente Médio e África	 Site <u>Dell.com</u> local específico do país ou <u>Dell.com/servicedescriptions</u>.* Além disso, os Clientes localizados na França, na Alemanha e no Reino Unido podem selecionar o URL aplicável abaixo: França: <u>Dell.fr/ConditionsGeneralesdeVente</u> Alemanha: <u>Dell.de/AGB</u> Reino Unido: <u>Dell.co.uk/terms</u> 	As Descrições de Serviço e outros documentos de serviço da Dell Technologies que você possa receber do vendedor não constituem um contrato entre você e a Dell Technologies, servindo apenas para descrever o conteúdo do Serviço que você está adquirindo desse vendedor, suas obrigações como destinatário do Serviço e os escopos e limitações de tal Serviço. Como consequência e neste contexto, qualquer referência a "Cliente" nesta Descrição de Serviço e em outros documentos de serviço da Dell Technologies será entendida como uma referência a você, enquanto as referências à Dell Technologies devem ser entendidas como uma referência à Dell Technologies como um provedor de serviços que está fornecendo o Serviço em nome do seu vendedor. Você não terá um relacionamento contratual direto com a Dell Technologies em relação ao Serviço descrito neste documento. Para evitar dúvidas, qualquer condição de pagamento ou contratual que seja por natureza diretamente relevante anenas
		entre um comprador e um vendedor não será aplicável a seu caso e seguirá o acordado entre você e seu vendedor.

* Os clientes podem acessar o site <u>Dell.com</u> local acessando <u>Dell.com</u> em um computador conectado à Internet em sua localidade ou escolhendo entre as opções do site "Choose a Region/Country" (Escolha uma Região/País) da Dell disponível em <u>Dell.com/content/public/choosecountry.aspx?c=us&l=en&s=gen</u>.

O cliente também concorda que, se renovar, modificar, estender ou continuar utilizando o Serviço além do período de vigência inicial, o Serviço estará sujeito à Descrição de Serviço em vigor na época e disponível para análise em <u>Dell.com/servicedescriptions</u>.

Ao fazer o pedido dos Serviços, receber a prestação dos Serviços, utilizar os Serviços ou o software associado ou clicar no botão/marcar a caixa "I Agree" (Eu concordo) ou similar em relação à sua compra no site Dell.com ou DellEMC.com, em uma interface da Internet ou um software da Dell Technologies, você concorda com o cumprimento desta Descrição de Serviço e com os acordos incorporados a ela por referência. Se você está firmando esta Descrição de Serviço em nome de uma empresa ou outra pessoa jurídica, você declara ter autoridade para vinculá-la a esta Descrição de Serviço e, nesse caso, "você" ou o "Cliente" deverá se subordinar à entidade em questão. Além de receberem esta Descrição de Serviço, os Clientes de determinados países também poderão ser solicitados a assinar um Formulário de pedido.



Termos e Condições Adicionais

Vigência do Serviço. Esta Descrição de Serviço se inicia na data registrada no Formulário de Pedido e continua até o final da vigência ("Vigência") indicada no Formulário de pedido. A quantidade de sistemas, licenças, instalações, implementações, endpoints ou usuários finais gerenciados para os quais o Cliente adquiriu um ou mais Serviços, o valor ou o preço e a Vigência aplicável a cada Serviço estão indicados no Formulário de Pedido do Cliente, conforme aplicável. Salvo se acordado por escrito de outra forma entre a Dell Technologies e o Cliente, as aquisições dos Serviços incluídos nesta Descrição de serviço serão exclusivamente para o uso interno do Cliente, e não para fins de revenda nem para agenciamento de serviços.

2. Informações adicionais importantes

- a) Reagendamento. Quaisquer alterações na programação após o agendamento deste Serviço devem ocorrer pelo menos 8 dias corridos antes da data agendada. Se o Cliente fizer um novo agendamento do serviço a sete dias ou menos da data agendada, será cobrada uma tarifa de novo agendamento, que não deverá exceder 25% do preço pago pelo Cliente para a prestação dos Serviços. Qualquer reagendamento do Serviço deverá ser confirmado pelo Cliente pelo menos 8 dias antes do início do Serviço.
- b) Pagamento do Hardware Adquirido com os Serviços. Salvo se acordado de outra forma por escrito, o pagamento do hardware, em nenhuma hipótese, dependerá da prestação ou do fornecimento dos Serviços comprados com ele.
- c) Limites do escopo dos Serviços. A Dell poderá se recusar a prestar o Serviço se, em sua opinião, a prestação do Serviço criar um risco não aceitável à Dell ou aos provedores de serviços da Dell ou se algum serviço solicitado estiver fora do escopo do Serviço. A Dell não se responsabiliza por nenhuma falha ou atraso na execução devido a causas fora de seu controle, inclusive a incapacidade de o Cliente cumprir com suas obrigações nos termos desta Descrição de serviço.
- d) Alterações no Escopo do Serviço. Salvo se acordado de outra forma por escrito com o Cliente, a Dell Technologies se reserva o direito de alterar o escopo dos Serviços em sessenta (60) dias antes do prévio aviso por escrito ao Cliente. Além das alterações causadas por editores e fabricantes de Produtos de Terceiros, o Cliente será notificado sobre qualquer alteração no escopo dos Serviços de Suporte da maneira declarada no Contrato atual entre a Dell Technologies e o Cliente.
- e) Privacidade. A Dell Technologies tratará todas as informações pessoais coletadas sob esta Descrição de serviço de acordo com a Declaração de Privacidade da Dell Technologies da jurisdição aplicável, todas as quais estão disponíveis em http://www.dell.com/localprivacy e cada uma delas é incorporada por referência.
- f) Serviços Opcionais. Serviços opcionais (inclusive suporte em pontos de necessidade, instalação, consultoria, serviços gerenciados, de suporte, profissionais ou de treinamento) podem estar disponíveis para compra com a Dell Technologies e variarão de acordo com a localização do Cliente. Os serviços opcionais podem exigir um contrato separado com a Dell Technologies. Na ausência de tal contrato, os serviços opcionais são fornecidos de acordo com esta Descrição de Serviço.
- g) Atribuição e Terceirização. A Dell Technologies pode terceirizar este Serviço e/ou atribuir esta Descrição de serviço a prestadores de serviço terceirizados qualificados, que prestarão o Serviço em nome da Dell Technologies.
- Cancelamento. A Dell Technologies pode cancelar este Serviço a qualquer momento durante a Vigência por qualquer um dos seguintes motivos:
 - a. O Cliente deixar de pagar o preço total ou parcial deste Serviço de acordo com os termos da fatura;
 - b. O Cliente for abusivo, ameaçador ou se recusar a cooperar com o analista assistente ou o técnico no local;
 - c. O Cliente deixar de respeitar os termos e condições estabelecidos nesta Descrição de serviço;
 - d. O Cliente solicita a substituição de componentes que excedem materialmente as taxas de falha padrão do componente e do sistema envolvidos, cujas taxas de falha são monitoradas constantemente. Consulte a seção de exclusão acima.

Se a Dell Technologies cancelar este Serviço de acordo com este parágrafo, ela enviará um aviso de cancelamento por escrito ao Cliente para o endereço indicado na fatura do Cliente. A notificação incluirá o motivo do cancelamento e a data de entrada em vigor do cancelamento, que não será em menos de 10 (dez) dias da data de envio do aviso de cancelamento pela Dell Technologies ao Cliente, a menos que alguma lei local exija outras provisões de cancelamento que não possam ser alteradas por contrato. Se a Dell Technologies cancelar este Serviço de acordo com este parágrafo, o Cliente não terá direito ao reembolso das tarifas pagas ou devidas à Dell Technologies.



i) Limitações Geográficas e Realocação.

- 1. A Dell Technologies pode não ser capaz de fornecer Serviços de Suporte a peças de reposição e no local em 4 horas com relação aos Equipamentos que estão fora da Área de Serviço da Dell Technologies. "Área da Dell Technologies Services" significa uma localização que está dentro de (i) um raio de 160 (cento e sessenta) quilômetros, ou 100 (cem) milhas, de carro de uma localização de serviço da Dell Technologies; e (ii) no mesmo país da localização de serviço da Dell Technologies, a menos que definido de outra forma no contrato vigente firmado com a Dell Technologies, hipótese em que a definição descrita no contrato em vigor prevalecerá. Para os Clientes situados na região EMEA, salvo indicação em contrário nesta Descrição de serviço ou no Contrato, o serviço no local está disponível até uma distância de 150 quilômetros da localização mais próxima de Logística da Dell Technologies (localização de retirada/entrega, ou PUDO). Antes da compra, entre em contato com seu representante de vendas para obter mais informações sobre a disponibilidade do serviço no local na região EMEA, de acordo com os locais de serviço da Dell Technologies.
- Este Serviço não está disponível em todos os locais. Se o Produto não estiver localizado na região geográfica que 2. corresponde ao local indicado nos registros de serviço da Dell Technologies para o Produto ou se os detalhes de configuração tiverem sido alterados e não comunicados à Dell Technologies, a Dell Technologies deverá primeiro requalificar seu Produto para o direito de suporte que você adquiriu antes de serem redefinidos os tempos de resposta aplicáveis para o Produto. As opções de Serviço, inclusive níveis de serviço, horas de suporte técnico e tempos de resposta no local, podem variar conforme a região geográfica e a configuração, e determinadas opções podem não estar disponíveis para compra na localização do Cliente. Portanto, entre em contato com seu representante de vendas para obter essas informações. A obrigação da Dell Technologies de prestação dos Serviços a Produtos realocados está sujeita a diversos fatores, incluindo, entre outros, disponibilidade local do Serviço, tarifas adicionais e inspeção e recertificação dos Produtos realocados, de acordo com as taxas atuais de consultoria de tempo e materiais da Dell Technologies. A menos que acordado de outro modo entre a Dell Technologies e o Cliente, nos casos em que as peças de serviço forem enviadas diretamente ao Cliente, o Cliente deverá ser capaz de receber as peças na localização de reparo dos Produtos. A Dell Technologies não será responsabilizada por atrasos no suporte devido à falha ou à recusa do Cliente em aceitar a remessa das peças. Os sistemas de armazenamento com vários componentes exigem contratos de opção de suporte ativos em todos os componentes de hardware e software do sistema a fim de receber todos os benefícios do contrato de suporte para a solução inteira. A menos que acordado de outra forma por escrito com o Cliente, a Dell Technologies reserva-se o direito de alterar o escopo dos Serviços de Suporte em sessenta (60) dias antes do prévio aviso por escrito ao Cliente.
- j) Ordem de Precedência. Salvo se acordado de outra forma por escrito entre as partes, se houver um conflito entre os termos de qualquer um dos documentos que compreendem este Contrato, os documentos prevalecerão na seguinte ordem: (i) esta Descrição de Serviço; (ii) o Contrato. As condições prevalecentes serão interpretadas na forma mais estrita possível para resolver o conflito preservando o maior número possível de condições não conflitantes, inclusive preservando as cláusulas não conflitantes no mesmo parágrafo, seção ou subseção.

ENTRE EM CONTATO CONOSCO

Para saber mais, entre em contato com o seu representante local ou revendedor autorizado.

Copyright © 2024 Dell Inc. ou suas subsidiárias. Todos os direitos reservados. Dell Technologies, Dell e outras marcas comerciais pertencem à Dell Inc. ou suas subsidiárias. Outras marcas comerciais podem ser marcas comerciais de seus respectivos proprietários. Publicado no Brasil.

A Dell Technologies acredita que as informações deste documento são precisas até a data da publicação. As informações estão sujeitas a alterações sem aviso prévio.



Dell Technologies Enterprise Systems Rail Sizing and Rack Compatibility Matrix

This document provides mounting features and key dimensions of the rack rails used for mounting many Dell Technologies enterprise systems and peripheral devices in a rack enclosure.



The information in this publication is provided "as is." Dell Inc. makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Use, copying, and distribution of any software described in this publication requires an applicable software license.

This document may contain certain words that are not consistent with Dell's current language guidelines. Dell plans to update the document over subsequent future releases to revise these words accordingly.

This document may contain language from third party content that is not under Dell's control and is not consistent with Dell's current guidelines for Dell's own content. When such third-party content is updated by the relevant third parties, this document will be revised accordingly.

Copyright © Jan 2025 | Version 5.0 Dell Inc. or its subsidiaries. All Rights Reserved. Dell Technologies, Dell, EMC, Dell EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be trademarks of their respective owners.

Contents

ntroduction	. 1
Considerations	. 1
Nounting interface	. 2
Rail types - System Installation Method	. 3
Cable Management Solutions	. 4
Backwards compatibility	. 5
Self-Adjusting Slide Feature	. 7
Definitions - Reference for Table 2	. 7

Figures

Figure 1.	Top view of right front EIA mounting flange	1
Figure 2.	System offset for round-hole racks	2
Figure 3.	Self-Adjusting Slide Feature	7

Tables

Table 1.	Dell Technologies server rails compatibility chart	. 5
Table 2.	DELL Technologies Rail Sizing Matrix	. 9
Table 3.	Dell Technologies rack compatibility matrix	22



Introduction

This document provides information about the mounting features and key dimensions of the rack rails used for mounting many Dell Technologies[™] enterprise systems and peripheral devices in a rack enclosure. This document also provides a compatibility summary for select Dell Technologies racks as well as some common third-party racks. Note that the product list is not all-inclusive and updates will be made as needed.

The dimensions provided in this document are for reference only. Some minor deviations due to manufacturing tolerances and variances should be expected.

Dell Technologies rail kits may not be compatible with racks from other vendors, however, all Dell Technologies rail kits are designed for compliance with all EIA-310-D and later revision specifications for 19-inch racks.

Considerations

Please pay attention to the footnotes indicated in the tables because they provide important information on using the rails in different racks and circumstances.

It is assumed that rack mount peripherals and cable bundles do not protrude into the space directly behind the systems.

Note that Dell Technologies rail kits with a Rail Identifier code have been designed to be compliant with the Server System Infrastructure (SSI) Specification for Computer Server Cabinet Enclosures & Racks, which specifies a minimum offset distance for return flanges on the rack mounting flanges to allow sufficient room for mounting the rail kits, as indicated in Figure 1. For more information about the Server System Infrastructure (SSI) Specification for Computer Server (SSI) Specification for Computer Server System Racks, see the SSI Forum at ssiforum.org.





Some third-party racks may not meet this requirement, and although Dell Technologies has made extensive efforts to accommodate as many third-party racks as possible, it is not feasible to provide a solution for every circumstance.

1 D<echnologies



Rack Types - 2-post and 4-post

Dell Technologies rail kits install into two different rack types with various flange hole designs. These rack types are broken down in Table 2 into 4-post and 2-post styles. 4-post rack types contain vertical mounting flanges with either square-hole, unthreaded round-hole, or threaded round-hole designs as part of the rack and rail interface. 2-post rack types generally contain threaded round-hole designs and require users to mount the server in either the front or center mount orientations. Only stab-in static rail designs that support 2-post rack configurations may be installed into a 2-post rack and commonly require additional hardware to secure the rails to the rack. For more information, refer to the definitions section for Table 2.

Mounting interface

The ReadyRails^M II mounting interface supports tool-less installation in 4-post square-hole and unthreaded round-hole racks as well as native support for tooled installation in threaded-hole racks. Note that installing this mounting interface in a square-hole rack allows the bracket to be placed flush against the mounting post, while installation in a round-hole rack results in a slight offset of approx. 6 mm from the mounting post, which also results in an approx. 6 mm bezel offset; refer to Figure 2.



The original **ReadyRails** mounting interface is used for both static and sliding rails, and it supports tool-less installation in 4-post square-hole and unthreaded round-hole racks. Static ReadyRails kits also support tooled installation in threaded-hole racks and 2-post racks. When installed in unthreaded round-hole racks, the original ReadyRails will also have the 6 mm offset from the mounting post that was discussed in the previous ReadyRails II paragraph. In order to install sliding ReadyRails kits into a threaded-hole rack, adapter brackets are required. 1U and 2U adapter bracket kits are available that support systems ranging from 1U to 5U in height.

The adapter bracket kits include six brackets to accommodate different rail lengths, plus four sets of custom screws in 10-32, 12-24, M5 and M6 thread sizes. The design of the brackets has been optimized to limit the forward shift of the system in the rack to only 17.3 mm. Depending on the depth of the rack used and the position of the mounting rails within the rack, it may be necessary to remove the system's bezel in order to close the front door of the rack. For the front door to close with the system



bezel installed, a minimum clearance of 58 mm is needed between the back surface of the door panel and the front face of the EIA flange.

The **RapidRails**[™] mounting interface supports tool-less installation in 4-post square-hole racks only, while the **VersaRails**[™] mounting interface supports tooled installation in 4-post square-hole and unthreaded round-hole racks. Mounting the VersaRails in threaded-hole racks is not recommended and is not supported by Dell Technologies.

The Generic mounting interface encompasses all other mounting interfaces outside of the ones listed above. Unless indicated to be tool-less, tools are required for installation.

Rail types - System Installation Method

Drop-in/Stab-in rails (Combo Rail) are a feature rich rail solution that allows a system to be fully extended out of the rack for service and the user has the option to install the system into the rail using a drop-in method like the ReadyRails sliding rails, or a stab-in method like the ReadyRails static rails. Drop-in/Stab-in rails support CMA or SRB applications. CMA and SRB applications must be detached in order to remove the inner member from the rails.

A "drop-in" design means that the system is installed vertically into the rails by inserting the standoffs on the sides of the system into the "J-slots" in the inner rail members with the rails in the fully extended position. The recommended method of installation is to first insert the rear standoffs on the system into the rear J-slots on the rails to free up a hand and then rotate the system down into the remaining J-slots while using the free hand to hold the rail against the side of the system.

A "stab-in" design means that the inner (chassis) rail members must first be attached to the sides of the system and then inserted into the outer (cabinet) members installed in the rack. For systems that are 2U and larger, it is recommended that two people perform this operation.

Sliding rails allow a system to be fully extended out of the rack for service. Most sliding rails support Cable Management Arms (CMAs) which enable the system to be extended out of the rack without disconnecting data/power cables at the rear of the system.

Unless otherwise indicated, all sliding rails are drop-in sliding rail design.

Static rails typically do not support the ability to service the system in the rack and are not compatible with the CMA. However, they do offer more flexibility in the types of racks and installations supported. Generally, there are two types of static rails: stab-in static and L-bracket static.

Stab-in static rails require the inner (chassis) rail members must first be attached to the sides of the system and then inserted into the outer (cabinet) members installed in the rack. For systems that are 2U and larger, it is recommended that two people perform this operation.

Stab-in Sliding rails require the inner (chassis) rail members must first be attached to the sides of the system and then inserted into the outer (cabinet) members installed in the rack and allow a system to be portion extended out of the rack for service. For systems that are 2U and larger, it is recommended that two people perform this operation. Most stab-in Sliding rails are compatible with CMA and SRB solutions.

L-bracket static rails do not support the ability to fully extend a system out of the rack into a service position. These rails typically are not compatible with cable management solutions unless otherwise indicated. Typically, equipment supported by L-bracket are customer serviceable from the front or rear of the rack.



Cable Management Solutions

To help manage the numerous cables associated with rack-mounted servers, a Cable Management Arm (CMA) or Strain Relief Bar (SRB) can be used. An optional CMA is offered with most sliding rails. CMAs attach on either the right or left side without tools.

Cable management arm (CMA) is a cable management accessory which connects to the rails behind the system. It allows a fully cabled system to be extended out of the rack into a service position.

Strain relief bar (SRB) is a cable management solution, which in most cases, attaches to the back of the rails via the strain relief bar brackets. Cables from the back of the chassis are placed across the top of the SRB and secured by straps.

SRBs are offered for select systems as an optional method for managing cables at the rear of the system due to the potential of a cable bundle size that exceeds the capacity of the CMA. The rail depth with a SRB is significantly less than that of a CMA, which in many cases, enables fitment of the rails in shallow racks. Cable service loops are required for systems on sliding rails to fully extend out of the rack for service.

Note that using a CMA or SRB with a deeper system may interfere with access to power distribution units (PDUs) in certain racks. If a configuration does not require CMA support, then the outer CMA mounting brackets can be removed from some of the sliding rail kits to reduce the overall length of the rails and eliminate potential interference with rear-mounted PDUs or the rack rear door.



Backwards compatibility

Some systems may offer backward compatibility with the rail kits from previous-generation systems. This is not always possible, because changes to chassis features, dimensions or weight can prevent older rail kits from being used with newer systems. Please refer to Table 1 for cross-generational compatibility of Dell Technologies servers and rails.

17 th Generation	Backwards compatibility with 16 th generation rails/CMAs											
product	Sliding rails	СМА	Static rails									
R670	✓	✓	1									
R770	✓	✓	1									
R570	✓	✓	1									
R470	X	X	X									
R7715	✓	✓	1									
R7725	✓	✓	1									
R6715	✓	✓	✓									
R6725	✓	✓	✓									

Table 1. Dell Technologies server rails compatibility chart

16 th Generation	Backwards con	Backwards compatibility with 15 th generation rails/CMAs											
product	Sliding rails	СМА	Static rails										
R260	N/A	N/A	✓										
R360	1	✓	✓										
R760	X	√*	X										
R660	✓	√*	✓										
R7615	X	X	X										
R7625	1	√*	✓										
R6615	X	X	X										
R6625	✓	√*	✓										
R760xd2	X	X	X										
R660xs	✓	√*	✓										
R760xs	✓	√*	✓										
R760xa	X	N/A	X										
R860	N/A	N/A	N/A										
R960	N/A	N/A	N/A										



15 th Generation	Backwards compatibility with 14 th generation rails/CMAs											
product	Sliding rails	СМА	Static rails									
R250	N/A	N/A	✓									
R350	✓	~	✓									
R650xs	✓	~	✓									
R450/R450xs	✓	✓	✓									
R750xs	X	✓	X									
R550/R550xs	X	✓	X									
R6515	✓	✓	✓									
R6525	N/A	N/A	N/A									
R650	X	✓	X									
R7515	✓	✓	✓									
R7525	X	X	X									
R750	✓	✓	✓									
R750xa	N/A	N/A	N/A									

14 th Generation	Backwards com	patibility with 13 th gene	eration rails/CMAs					
product	Sliding rails	СМА	Static rails					
R240	N/A	N/A	✓					
R340	X	✓	✓					
R440	X	✓	✓					
R540/R540xd	✓	✓	✓					
R640	✓	✓	✓					
R740/R740xd	✓	✓	✓					
R740xd2	N/A	N/A	N/A					
R840/940xa	X	X	X					
R940	X	✓	N/A					
C4140	N/A	N/A	✓					
C64xx	N/A	N/A	✓					
T440	Х	~	N/A					
T640	✓	✓	N/A					

Compatible
 X - Not compatible
 *Only with the previous generation sliding rail

6 D≪LLTechnologies



Self-Adjusting Slide Feature

For many 1U and 2U systems, rails have been standardized with a slim design that holds a wide system chassis to accommodate more features and functions. They also have a self-adjusting slide feature that accommodates different depths of systems, offering compatibility across multiple platform models. Refer to Figure 3 for an illustration of how the self-adjusting slide feature works.



The rail adjustability range when the rails are installed in a rack is the same regardless of system depth since the feature is not utilized until a system is installed. If the system being installed in the rails requires this feature, the minimum rail adjustability limit is increased by the amount of travel the slide body needs to slide back to support the system. The minimum rail adjustability limit is documented in the resources listed at the end of this notice.

Users who have systems that utilize the feature might observe a slight amount of additional resistance from a spring in each rail when the system is almost completely installed in the rack. For most rails, the instance when the resistance is observed is within the final 55 mm of translation before the slam latch is engaged with the rail.

The rail slide-adjusting feature can be found on both sliding and drop-in/stab-in rail types. The rail adjustability range (mm) values listed in Table 2 for products that utilize this rail feature have been flagged with a footnote.

Definitions - Reference for Table 2

Rail identifier is a two-character code used on most rail kits to indicate compatibility between rails and systems. The twocharacter code consists of a letter followed by a one or two-digit number. It is typically located on a front inside surface on both the left and right sliding rail and drop-in/stab-in rail members. If there is a component of the rail kit that is attached to the chassis prior to installing the system into a rack, such as with the stab-in static rails, the identifier is located closer to the center of the component.

Square-hole describes a 4-post rack mounting flange type where rails utilize Square holes sized according to EIA-310-D standard for mounting.

Round-hole describes a 4-post rack mounting flange type where rails utilize Unthreaded-Round holes sized according to EIA-310-D standard for mounting.



Threaded-hole describes a 4-post rack mounting flange type where rails utilize Threaded-round holes for mounting. Threaded-round holes may require additional hardware for mounting and hardware may vary by thread type. See footnotes in table 2 for specific information on threaded-round hole mounting.

Mounting interface describes the type of rail bracket design used for mounting the rail in the rack.

Rail adjustability range represents the allowable distance between the outside-facing surfaces of the front and rear mounting posts of the rack when a system is fully installed. This does not include the portion of the rail kit or other rail components that may extend beyond the mounting posts.

Rail depth represents the minimum depth of the rail as measured from the rack front mounting posts when the rail rear bracket is positioned all the way forward. The rail may extend beyond the rear bracket, particularly for sliding rail kits to support CMA or SRB attachment. In some instances, the chassis may extend beyond the minimum rail depth, and in such cases, please refer to the footnotes in Table 2.

Table 2. DELL Technologies Rail Sizing Matrix

			Pail	Mounting		Rack types supported					Rail adjustability range (mm)							Rail depth (mm)	
		Product	Rail identifier	Mounting interface	Rail type		4-Post		2-	Post	Squ	are	Ro	und	Threaded		without	with	
						Square	Round	Thread	Flush	Center	Min	Max	Min	Max	Min	Max	CMA/SRB	CMA(SRB)	
			Α7	ReadyRails II	Sliding	~	✓	√ a,c,d	X	x	631	868	617	861	631	883	720 ^b	845	
		R640 (8-HDD)	A8	ReadyRails	Stab-in Static	~	~	√ a,c	√ a,c	✓c	608	879	594	872	610	898	622 ⁿ	-	
			A10	Generic Tool-less	Drop-in/ Stab-in	~	~	✓ ₩	x	x	559	931	559	931	559	931	720 ^b	845	
			Α7	ReadyRails II	Sliding	~	×	√ a,c,d	X	x	681 ^p	868	667 ^p	861	681 ^p	883	770 ^b	895	
		R640 (4-HDD/10-HDD)	A8	ReadyRails	Stab-in Static	~	~	√ a,c	√ a,c	√c	608	879	594	872	610	898	622 ⁿ	-	
			A10	Generic Tool-less	Drop-in/ Stab-in	~	×	✓ ₩	X	x	613 ^p	931	613 ^p	931	613 ^p	931	770 ^b	895	
/RS	:dge TM	R6525 (8-HDD) R650 (8-HDD) R660/R6615/R6625	A15	ReadyRails II	Sliding	~	×	√ a,c,d	X	x	631	868	617	862	631	884	736 ^b	862 (770/792)	
SER	owerE		A14	ReadyRails	Stab-in Static	~	×	√ a,c	√ a,c	✓c	608	880	594	870	605	893	622 ⁿ	-	
	đ		A16	Generic Tool-less	Drop-in/ Stab-in	*	~	✓ ₩	X	x	559	994	559	994	559	944	736 ^b	862 (770/792)	
		R6525 (4-HDD/10-HDD)	A15	ReadyRails II	Sliding	~	×	√ a,c,d	X	x	631	868	617	862	631	884	787 ^ь	913 (821/843)	
		R650 (4-HDD/10-HDD) R660/ R6615/R6625	A14	ReadyRails	Stab-in Static	~	~	√ a,c	√ a,c	√c	608	880	594	870	605	893	622 ⁿ	-	
		(4-HDD/10-HDD/E3/LO-2) R670/R6725/R6715/R470	A16	Generic Tool-less	Drop-in/ Stab-in	~	~	1	x	X	610	994	610	994	610	994	787 ^ь	913 (821/843)	
		R340 (8-HDD)	A12	ReadyRails II	Sliding	~	~	√ a,c,d	X	X	631	868	617	861	631	883	720 ^b	845	
	R350 (8-HDD) R360 (8-HDD)	A8	ReadyRails	Stab-in Static	~	✓	√ a,c	√ a,c	✓c	608	879	594	872	610	898	622 ⁿ	-		



					Rack types sup			ported			Rail a	Rail depth (mm)					
	Product	Rail identifier	Mounting interface	Rail type		4-Post		2-	Post	Squ	are	Ro	und	Threaded		without	with
					Square	Round	Thread	Flush	Center	Min	Max	Min	Max	Min	Max	CMA/SRB	CMA(SRB)
	R340 (4-HDD)	A12	ReadyRails II	Sliding	~	×	√ a,c,d	X	X	681 ₽	868	667 ^p	861	681 ^p	883	770 ^b	895
	R360 (4-HDD)	A8	ReadyRails	Stab-in Static	~	~	√ a,c	√ a,c	√c	608	879	594	872	610	898	622 ⁿ	-
	R440 (8-HDD) R450 (8-HDD) R6415 (8-HDD)	A8	ReadyRails	Stab-in Static	~	~	√ a,c	√ a,c	√ c	608	879	594	872	610	898	622 ⁿ	-
	R6515 (8-HDD) R650xs (0-HDD/8-HDD) R660xs(0-HDD/8-HDD)	A11	Generic Tool-less	Drop-in/ Stab-in	~	~	✓ ₩	x	x	559	931	559	931	559	931	720 ^b	845 (761/783)
	R440 (4-HDD/10-HDD) R450 (4-HDD) R6415 (4-HDD/10-HDD) R6515 (4-HDD/10-HDD) R650xs (4-HDD/10-HDD) R650xs (8-HDD NVME) R660xs (4-HDD/10-HDD)	A8	ReadyRails	Stab-in Static	~	~	√a,c	√ a,c	√c	608	879	594	872	610	898	622 ⁿ	-
		A11	Generic Tool-less	Drop-in/ Stab-in	*	~	*	X	X	609 ^p	931	609 ^p	931	609 ^p	931	770 ^b	895 (811/833)
		B6	ReadyRails II	Sliding	~	~	√ a,c,d	X	X	631	868	617	861	631	883	714 ^b	845
	R540/R540xd R740/R740xd/ R7415/R7425/R7515	B4	ReadyRails	Stab-in Static	~	✓	√ a,c	√ a,c	√c	608	879	594	872	610	898	622 ⁿ	-
	K/415/K/425/K/515	B13	Generic Tool-less	Drop-in/ Stab-in	~	~	✓ ₩	X	x	559	931	559	931	559	931	714 ^b	845
		B20	ReadyRails	Stab-in Static	~	~	√ a,c	√ a,c	√c	608	879	594	872	610	898	622 ⁿ	-
	R550 R750xs	B21	ReadyRails II	Sliding	~	~	√ a,c,d	X	X	631	868	617	861	631	883	714 ^b	845 (749/769)
	R760xs	B22	Generic Tool-less	Drop-in/ Stab-in	~	~	~	x	x	559	931	559	931	559	931	714 ^b	845 (749/769)
	87565	B6	ReadyRails II	Sliding	~	~	√ a,c,d	X	X	685 ^p	868	671 ^p	861	685 ^p	883	766 ^b	895 (802/822)
	R7525 R750	B4	ReadyRails	Stab-in Static	~	~	√ a,c	√ a,c	√c	608	879	594	872	608	898	622 ⁿ	-
		B13	Generic Tool-less	Drop-in/ Stab-in	~	×	✓ ₩	X	X	609 ^p	931	609 ^p	931	609 ^p	931	779 ^ь	899 (802/822)

D&LLTechnologies

Inserido ao protocolo 22.951.206-4 por: Pedro Henrique Golin Linhares em: 31/03/2025 12:49. A autenticidade deste documento pode ser validada no endereço: https://www.eprotocolo.pr.gov.br/spiweb/validarDocumento com o código: 692a5ef82beb173ed33dad4ee6505546.





11

OTOCO

Fls. <u>362</u> Mov. <u>32</u>

ADO DO

D&LLTechnologies

						Rack types supp		ported			Rail a	Rail depth (mm)					
	Product	Rail identifier	Mounting interface	Rail type		4-Post		2-	Post	Squ	are	Round		Threaded		without	with
					Square	Round	Thread	Flush	Center	Min	Max	Min	Max	Min	Max	CMA/SRB	CMA(SRB)
	EV2/EV2c	B10	ReadyRails II	Sliding	~	×	√ a,c,d	x	x	677	815	665	809	677	830	836	888
	FAZ/FAZ3	B11	ReadyRails II	Stab-in Static	~	~	√ a,c	X	X	644	916	632	910	644	930	828	-
	C4130/C4140	Α9	ReadyRails II	Stab-in Static ^h	*	✓	√ a,c,d	X	X	643	916	631	910	643	930	766	-
	T640	C4	ReadyRails II	Sliding	1	×	√ a,c,d	x	x	686	756	672	749	686	771	756	840
	T440	C2	ReadyRails II	Sliding	~	×	√ a,c,d	X	x	686	756	672	749	686	771	760	840
	VRTX	С3	ReadyRails II	Sliding	~	×	√ a,c,d	X	x	608	915	594	908	608	930	756	845
	R240/R250/R260	A4	ReadyRails	Stab-in Static	~	×	√ a,c	√ a,c	√c	608	879	594	872	610	898	622 ⁿ	-
	M1000e	-	RapidRails	L-Bracket Static	*	x	X	X	x	712	755	-	-	-	-	703	-
		-	VersaRails	L-Bracket Static	~	~	X	X	X	706	755	706	755	-	-	703	-
	MX7000	C5	ReadyRails II	L-Bracket Static	~	×	x	X	x	592	876	578	869	-	-	m	(901)
		A12	ReadyRails II	Sliding	~	×	√ a,c,d	X	x	681 ^p	868	667 ^p	861	681 ^p	883	770 ^b	895
	XR2	Α4	ReadyRails	Stab-in Static	~	~	√ a,c	√a,c	✓c	608	879	594	872	610	898	622 ⁿ	-
		-	Generic	Stab-in Static ^t	√ g	√g	√g	X	x	464	766	464	766	464	766	464 ⁿ	-
	VR44/VR42	A20	Generic Tool-less	Stab-in Sliding	✓	×	×	√a	√a	472	757	472	757	472	757	445 ^b	605 (498/520)
	XK11/XK12	A21 ^s	Generic Tool-less	Stab-in Sliding	X	×	x	X	x	-	-	458	589	-	-	-	-
e C	C1100	-	Generic Tool-less	Sliding	~	~	x	x	X	665	950	665	950	-	-	685	-
erEdg	C2100	-	Generic	Sliding	~	✓	~	X	x	664	1110	664	1110	664	1110	720	-
Pow	C410x	-	VersaRails	Stab-in Static	~	✓	x	X	x	737	972	737	972	-	-	734	-

12

OTOC

Fls. <u>363</u> Mov. <u>32</u>

D&LLTechnologies

					- 1	Rack t	ypes sup	ported			Rail a	ljustabi	lity rang	e (mm)		Rail dep	th (mm)
	Product	Rail identifier	Mounting interface	Rail type		4-Post		2-	Post	Squ	are	Ro	und	Thre	aded	without	with
					Square	Round	Thread	Flush	Center	Min	Max	Min	Max	Min	Max	CMA/SRB	CMA(SRB)
	C5xxx	-	Generic Tool-less	L-Bracket Static	×	×	X	X	X	708	947	708	947	-	-	705	-
	C64xx C65xx C66xx	-	Generic Tool-less	L-Bracket Static	~	~	X	X	X	609ª	917	609ª	917	-	-	-	-
	C8000	-	Generic Tool-less	L-Bracket Static	✓	~	X	x	X	708	946	708	946	-	-	713	-
		B20	ReadyRails	Stab-in Static	~	~	√ a,c	√ a,c	√c	608	879	594	872	610	898	622 ⁿ	-
	HS5620	B21	ReadyRails II	Sliding	*	~	√ a,c,d	x	X	631	868	617	861	631	883	714 ^b	845 (749/769)
		B22	Generic Tool-less	Drop-in/ Stab-in	~	~	✓	x	X	559	931	559	931	559	931	714 ^b	845 (749/769)
SP	HS5610 Cold Aisle	A22	ReadyRails II	Stab-in, Static	*	~	√ a,c,d	X	X	608	916	594	909	610	641	897	-
Edge C		A8	ReadyRails	Stab-in Static	*	~	√ a,c	√ a,c	√c	608	879	594	872	610	898	622 ⁿ	-
Power	(סמח-פיססח-ס)סו פכנח	A11	Generic Tool-less	Drop-in/ Stab-in	~	~	✓ ₩	X	X	559	931	559	931	559	931	720 ^b	845 (761/783)
		A8	ReadyRails	Stab-in Static	*	~	√ a,c	√ a,c	√c	608	879	594	872	610	898	622 ⁿ	-
	(עשריטו ששווידי)טו טכנוו	A11	Generic Tool-less	Drop-in/ Stab-in	~	*	√ ₩	x	x	609 ^p	931	609 ^p	931	609 ^p	931	770 ^b	895 (811/833)
XR		A23			~	~	√a	X	X	552	763	552	763	552	763	559 ^b	
erEdge	XR4000R	A24	Generic Tool-less	Stab-in Sliding	✓	×	√a	X	X	342	554	342	554	342	554	358 ^b	-
Powe		A25 ^s		Ĩ	x	~	х	x	x	-	-	426	569	-	-	-	

D&LLTechnologies

Inserido ao protocolo 22.951.206-4 por: Pedro Henrique Golin Linhares em: 31/03/2025 12:49. A autenticidade deste documento pode ser validada no endereço: https://www.eprotocolo.pr.gov.br/spiweb/validarDocumento com o código: 692a5ef82beb173ed33dad4ee6505546.



						Rack t	ypes sup	ported			Rail ad	ljustabi	lity rang	e (mm)		Rail dep	th (mm)
	Product	Rail identifier	Mounting interface	Rail type		4-Post		2-	Post	Squ	are	Ro	und	Thre	aded	without	with
					Square	Round	Thread	Flush	Center	Min	Max	Min	Max	Min	Max	CMA/SRB	CMA(SRB)
	VEROOR	B31	Generic Tool-less	L-Bracket Static	~	~	√a	~	*	293	516	293	516	293	516	448	-
	AROUUUR	B36	Generic Tool-less	L-Bracket Static	~	~	√a	~	~	293	516	293	516	293	516	448	
	VD5/40	A26	Generic	Stab-in	~	~	√a	x	x	472	766	472	766	472	766	544	672 (581/604)
	XK3610	A27 ^s	Tool-less	Sliding	x	~	x	x	x	-	-	432	562	-	-	-	-
	VR7/00	B29	Generic	Stab-in	~	~	√a	~	~	472	794	472	794	472	794	544	658 (580/603)
	XR7620	B30 ^s	Tool-less	Sliding	x	~	x	x	x	-	-	432	562	-	-	-	-
	XE2420	_	Generic	Stab-in Static Standard	~	~	√a,r	X	x	625	883	625	883	625	891	645 ⁿ	_
	AL2420		benerie	Stab-in Static Short	~	~	√ a,r	x	x	395	547	395	547	395	554	445 ⁿ	
	XE8640	B27	Generic Tool-les	Stab-in Sliding	~	×	√ ₩	x	x	609	924	609	924	609	924	827	935
lge XE	XE9680	B28	Generic Tool-les	Stab-in Sliding	~	~	√w	X	X	607	924	607	924	607	924	961	1075
owerE	XE9640	B25	Generic Tool-less	Stab-in Sliding	~	~	✓w	x	X	606	920	606	920	606	920	912	(945/967)
	XE9680L/XE9685L	B38	Generic Tool-less	Stab-in Sliding	~	~	√ ₩	X	x	608	924	608	924	608	924	970	-
	XE7740/XE7745	B37	Generic Tool-less	Stab-in Sliding	~	~	√w	x	x	608	924	608	924	608	924	832	-

D&LLTechnologies

Inserido ao protocolo 22.951.206-4 por: Pedro Henrique Golin Linhares em: 31/03/2025 12:49. A autenticidade deste documento pode ser validada no endereço: https://www.eprotocolo.pr.gov.br/spiweb/validarDocumento com o código: 692a5ef82beb173ed33dad4ee6505546.



							Rack t	ypes sup	ported			Rail ac	ljustabi	lity rang	e (mm)		Rail dep	th (mm)
		Product	Rail identifier	Mounting interface	Rail type		4-Post		2-	Post	Squ	are	Ro	und	Thre	aded	without	with
						Square	Round	Thread	Flush	Center	Min	Max	Min	Max	Min	Max	CMA/SRB	CMA(SRB)
		1081AD/2161AD 1082DS/2162DS 4322DS	Α5	ReadyRails	Stab-in Static	~	~	✓	~	~	496	770	482	763	488	794	506 ^Q	-
	KVM	180AS/2160AS 2161DS/2161DS-2 4161DS	-	Generic	Stab-in Static	>	~	*	~	x	686	737	686	737	686	737	686	-
		2321DS	-	Generic	Stab-in Static	~	~	1	~	x	533	737	533	737	533	737	533	-
		PC8132/PC8132F PC8164/PC8164F	A5	ReadyRails	Stab-in Static	~	*	~	~	✓	496	770	482	763	488	794	506 ^Q	-
		S4820T/S6000	A5	ReadyRails	Stab-in Static	×	~	✓	~	~	496	770	482	763	488	794	506 ^Q	-
		S5000	-	Generic	Stab-in Static	~	~	~	X	x	680	830	680	830	680	830	680	-
HES		Z9100	A5	ReadyRail	Stab-in Static	~	×	~	✓	~	496	770	482	763	488	794	506 ^Q	-
WITC		S4248	A5	ReadyRail	Stab-in Static	~	×	~	~	~	496	770	482	763	488	794	506 ^Q	-
01	ing	S41xx	A5	ReadyRail	Stab-in Static	~	~	~	~	~	496	770	482	763	488	764	506 ^Q	-
	twork	S4048/S4048T	A5	ReadyRail	Stab-in Static	~	~	~	~	~	496	770	482	763	488	764	506 ^Q	-
	Ne	S6010	A5	ReadyRail	Stab-in Static	~	~	~	~	~	496	770	482	763	488	764	506 ^Q	-
		S3048	A5	ReadyRail	Stab-in Static	×	~	~	~	~	496	770	482	763	488	764	506 ^Q	-
		S6100	B9	ReadyRails II	L-Bracket Static	~	~	√ a,c,d	X	x	595	914	581	907	595	929	600	-
		S6100NEBS	-	Generic	Stab-in Static	X	X	x	~	x	-	-	-	-	-	-	-	-
		N2128PX-ON	-	Generic	Stab-in Static	X	X	x	✓	x	-	-	-	-	-	-	-	-
		N3132PX-ON	A5	ReadyRails	Stab-in Static	~	✓	~	✓	~	496	770	482	763	488	764	506 ^Q	-



D&LLTechnologies

Inserido ao protocolo 22.951.206-4 por: Pedro Henrique Golin Linhares em: 31/03/2025 12:49. A autenticidade deste documento pode ser validada no endereço: https://www.eprotocolo.pr.gov.br/spiweb/validarDocumento com o código: 692a5ef82beb173ed33dad4ee6505546.

						Rack t	ypes sup	ported			Rail a	djustabi	ility rang	e (mm)		Rail dep	th (mm)
	Product	Rail identifier	Mounting interface	Rail type		4-Post		2-	Post	Squ	are	Ro	und	Thre	aded	without	with
					Square	Round	Thread	Flush	Center	Min	Max	Min	Max	Min	Max	CMA/SRB	CMA(SRB)
	N1108T/N1108P	-	Generic	Stab-in Static	x	X	x	~	x	-	-	-	-	-	-	-	-
	N1124T/N1124P	-	Generic	Stab-in Static	X	X	X	~	X	-	-	-	-	-	-	-	-
	N1148T/N1148P	-	Generic	Stab-in Static	X	X	X	~	X	-	-	-	-	-	-	-	-
	N3024/N3048	A5	ReadyRails	Stab-in Static	*	*	*	~	*	496	770	482	763	488	764	506 ^Q	-
	\$5148	A5	ReadyRails	Stab-in Static	*	×	*	~	×	496	770	482	763	488	764	506 ^Q	-
	S31xx	A5	ReadyRail	Stab-in Static	*	×	*	~	*	496	770	482	763	488	764	506 ^Q	-
	N30xx	A5	ReadyRail	Stab-in Static	*	×	*	~	*	496	770	482	763	488	764	506 ^Q	-
	P7010	B6	ReadyRails II	Sliding	*	×	√ a,c,d	X	X	631	868	617	861	631	883	714 ^b	845
	67910	B4	ReadyRails	Stab-in Static	*	×	√ a,c	√a,c	√c	608	879	594	872	610	898	622	-
	Precision 3930 Rack	Α4	ReadyRails	Stab-in Static	~	×	√ a,c	√ a,c	✓c	608	879	594	872	610	898	622 ⁿ	-
10		B6	ReadyRails II	Sliding	×	×	√ a,c,d	X	x	631	868	617	861	631	883	714 ^b	845
	Precision 7920 Rack	B4	ReadyRails	Stab-in Static	~	×	√a,c	√ a,c	√c	608	879	594	872	610	898	622 ⁿ	-
RKST⊿		B13	Generic Tool-less	Drop-in/ Stab-in	*	×	✓ ₩	X	X	607 ^p	931	607 ^p	931	607 ^p	931	714 ^b	845
MO		B20	ReadyRails	Stab-in Static	~	×	√ a,c	√ a,c	✓c	608	879	594	872	610	898	622 ⁿ	-
	Procision 7960 Pack	B21	ReadyRails II	Sliding	~	✓	√ a,c,d	X	X	685 ^p	868	671 ^p	861	685 ^p	883	766 ^b	895 (802/822)
	FIECISION / YOU RACK	B22	Generic Tool-less	Drop-in/ Stab-in	~	~	~	x	X	609 ^p	931	609 ^p	931	609 ^p	931	766 ^b	895 (802/822)
	T7600/T7610	C2	ReadyRails II	Sliding	~	~	√a,c,d	X	X	686	756	672	749	686	771	760	840

16

OTOC

Fls. <u>367</u> Mov. <u>32</u>

D&LLTechnologies

						Compa	Rack t	ypes sup	ported			Rail ad	ljustabi	lity rang	e (mm)		Rail dep	th (mm)
		Product	Rail identifier	Mounting interface	Rail type		4-Post		2-	Post	Squ	are	Ro	und	Thre	aded	without	with
						Square	Round	Thread	Flush	Center	Min	Max	Min	Max	Min	Max	CMA/SRB	CMA(SRB)
		R5500/R7610	B2	ReadyRails	Sliding	~	×	√f	x	x	686	883	672	876	651	897	755 ⁵	883
		FPM185 (without KVM)	-	ReadyRails II	Sliding	~	✓	√ a,c,d	x	X	604	900	590	893	604	914	-	611
	¥	FPM185 (with KVM)	-	ReadyRails II	Sliding	✓	~	√ a,c,d	X	X	705	900	691	893	705	914	-	715
:	2	1750	-	RapidRails	Sliding	~	x	X	X	X	714	755	-	-	-	-	-	787
		1717	-	VersaRails	Sliding	~	~	X	X	X	709	755	709	755	-	-	-	787
	SAU	Dell Rack Mount UPS Family	B5	ReadyRails	Stab-in Static	~	~	√f	х	x	518	769	504	762	483	783	526	-
	OTHER	1U Fixed Equipment Shelf	A4	ReadyRails	Stab-in Static	*	*	√ a,c	√ a,c	√ c	608	879	594	872	610	898	622	-
	PowerMax	PROXPE1 PROXPE2	B14	-	L-Bracket Static	√u	✓u	√u,v	x	x	558	914	558	914	558	914	600	1015
			Α7	ReadyRails II	Sliding	~	×	√ a,c,d	х	x	631	868	617	861	631	883	720 ^b	845
		NX3300/NX400	A8	ReadyRails	Stab-in Static	~	~	√ a,c	√a,c	✓c	608	879	594	872	610	898	622	-
RAGE	×	NX2200	B6	ReadyRails II	Sliding	*	*	√ a,c,d	X	x	631	868	617	861	631	883	714 ^b	845
STO	Vault	NA3200	B4	ReadyRails	Stab-in Static	~	×	√ a,c	√a,c	√c	608	879	594	872	610	898	622	-
	ower	NV2E00 Controller	Α3	ReadyRails	Sliding	*	~	√ e	X	x	686	883	672	876	651	897	714 ^b	835
		NASSOU COntroller	A4	ReadyRails	Stab-in Static	~	✓	√ a,c	√ a,c	√c	608	879	594	872	610	898	622	-
		NX3500 UPS	A4	ReadyRails	Stab-in Static	~	~	√ a,c	√ a,c	√c	608	879	594	872	610	898	622	-
		DX6000G	A4	ReadyRails	Stab-in Static	~	~	√ a,c	√ a,c	√c	608	879	594	872	610	898	622	-

17

D&LLTechnologies



					Rack t	ypes sup	ported			Rail a	djustabi	lity rang	e (mm)		Rail dep	th (mm)
Product	Rail identifier	Mounting interface	Rail type		4-Post		2-	Post	Squ	are	Ro	und	Thre	aded	without	with
				Square	Round	Thread	Flush	Center	Min	Max	Min	Max	Min	Max	CMA/SRB	CMA(SRB)
	A6	ReadyRails	Stab-in Static	~	✓	√a,c	√ a,c	√c	508 ^c	751	494 ^c	744	519 ^c	762	515 ^c 376 ^d	-
NY 300/DY 600/S	A3	ReadyRails	Sliding	×	✓	√e	X	x	686	883	672	876	651	897	714 ^b	835
NA300/DA00043	A4	ReadyRails	Stab-in Static	~	×	√ a,c	√ a,c	√c	608	879	594	872	610	898	622	-
NY 2000/DY 6000	B1	ReadyRails	Sliding	*	~	√f	X	x	692	756	678	749	657	770	751	840
NX3000/DX8000	A2	ReadyRails	Stab-in Static	*	~	√ a,c	√ a,c	√c	588	828	574	821	592	846	608	-
NX3100/DL2200	B3	ReadyRails	Sliding	*	×	√f	X	X	686	883	672	876	651	897	714 ^b	845
DX6012S/DR4000	B4	ReadyRails	Stab-in Static	*	~	√ a,c	√a,c	√c	608	879	594	872	610	898	622	-
MD3060e/MD3660	-	VersaRail	L-Bracket Static	*	~	x	x	x	611	791	611	791	-	-	620	-
	B9	ReadyRails II	L-Bracket Static	*	×	√ a,c,d	x	x	595	914	581	907	595	929	600	-
MD12xx/14xx/32xx/ 36xx/NX36xx	-	RapidRails	L-Bracket Static	*	X	x	X	x	732	758	-	-	-	-	729	-
	-	VersaRails	L-Bracket Static	*	✓	X	X	X	714	758	714	758	-	-	721	-
MD1120	-	RapidRails	L-Bracket Static	*	X	x	X	x	732	759	-	-	-	-	729	-
MD 1120	-	VersaRails	L-Bracket Static	*	✓	X	X	X	714	759	714	759	-	-	721	-
MD1000/MD2000	-	RapidRails	L-Bracket Static	*	x	x	x	x	732	758	-	-	-	-	735	-
MD 1000/MD 3000	-	VersaRails	L-Bracket Static	~	×	X	X	X	714	758	714	758	-	-	735	-
	B7	ReadyRails	Stab-in Static	~	✓	√ a,c	√ a,c	√c	588	828	574	821	592	846	608	-
	-	RapidRails	Sliding	~	X	X	X	X	722	750	-	-	-	-	792	870
F V I I 7 I / F V I I 7A	-	VersaRails	Sliding	~	✓	x	X	X	701	745	701	745	-	-	792	870

18

D&LLTechnologies



	Dett recimotogico Ent					Rack t	ypes sup	ported			Rail a	ljustabi	lity rang	e (mm)		Rail dep	th (mm)
	Product	Rail identifier	Mounting interface	Rail type		4-Post		2-	Post	Squ	are	Ro	und	Thre	aded	without	with
					Square	Round	Thread	Flush	Center	Min	Max	Min	Max	Min	Max	CMA/SRB	CMA(SRB)
	D)/4.2.4T	-	RapidRails	L-Bracket Static	~	x	x	x	X	729	755	-	-	-	-	732	-
	PV1241	-	VersaRails	L-Bracket Static	~	~	x	X	x	711	755	711	755	-	-	732	-
		A1	ReadyRails	Sliding	~	~	√e	X	x	692	756	678	749	657	770	768 ^b	887
	FS7500 Controller	A2	ReadyRails	Stab-in Static	~	~	√ a,c	√a,c	✓c	588	828	574	821	592	846	608	-
-	FS7500 UPS	Α4	ReadyRails	Stab-in Static	~	1	√ a,c	√a,c	√c	608	879	594	872	610	898	622	-
-ogic ["]		B9	ReadyRails II	L-Bracket Static	~	1	√ a,c,d	X	X	595	914	581	907	595	929	600	-
Equall	FS76xx/PS41xx PS61xx	-	RapidRails	L-Bracket Static	~	X	x	X	x	732	758	-	-	-	-	729	-
_		-	VersaRails	L-Bracket Static	~	1	x	X	x	714	758	714	758	-	-	721	-
	PS6500/6510	-	ReadyRails	Sliding	~	1	√ a,c	X	X	597	793	583	786	605	817	885	885
	PS4000/6000/6010	-	Generic	L-Bracket Static	~	√a	√ a	X	X	616	914	616	914	616	914	616	-
	568000	B6	ReadyRails II	Sliding	~	×	√ a,c,d	X	X	631	868	617	861	631	883	714 ^b	845
	308000	B4	ReadyRails	Stab-in Static	~	1	√ a,c	√a,c	✓c	608	879	594	872	610	898	622	-
ent™	SC20xx/SC40xx	-	Generic	L-Bracket Static	~	~	√ A	X	X	611	914	614	914	614	914	-	-
mpell	SC20xx/SC40xx	B9	ReadyRails II	L-Bracket Static	~	1	√ a,c,d	X	X	595	914	581	907	595	929	600	-
ell Co	SC2xx/FS86xx	-	RapidRails	L-Bracket Static	~	X	X	X	X	732	758	-	-	-	-	729	-
Δ		-	VersaRails	L-Bracket Static	~	~	X	X	X	714	758	714	758	-	-	721	-
	SCV30xx SC50xx SC7020	В9	ReadyRails II	L-Bracket Static	~	~	√ a,c,d	X	x	595	914	581	907	595	929	600	-

D&LLTechnologies





						Rack t	ypes sup	ported			Rail a	djustabi	lity rang	e (mm)		Rail dep	th (mm)
	Product	Rail identifier	Mounting interface	Rail type		4-Post		2-	Post	Squ	are	Ro	und	Thre	aded	without	with
					Square	Round	Thread	Flush	Center	Min	Max	Min	Max	Min	Max	CMA/SRB	CMA(SRB)
	Series 40	-	Generic	Sliding	~	√g	√ g	x	x	669	923	669	923	707 ^g	961 ^g	693	-
	Fibre Channel	-	Generic	Stab-in Static ^h	~	×	*	X	x	606	910	606	910	606	910	598	-
	SAS (new rails)	-	Generic	Stab-in Static ^h	~	×	X	X	X	606	910	606	910	606	910	598	-
	SAS (old rails)	-	Generic	Stab-in Static ^h	~	~	1	X	X	682	885	682	885	682	885	598	-
	NAS Gen3	-	Generic	Sliding	√i	✓i	√i	X	X	652	854	652	854	652	854	810	-

20

D&LLTechnologies

Notes:

- ^a Minor conversion required
- ^b With CMA brackets removed
- ^c Mounting screws not included in the kit
- ^d Mounting screw head diameter must be 10 mm or less
- ^e Requires the 1U Threaded Rack Adapter Brackets Kit (Dell PN 8Y19G), which shifts the system forward in the rack by 17.3 mm
- ^f Requires the 2U Threaded Rack Adapter Brackets Kit (Dell PN PKCR1), which shifts the system forward in the rack by 17.3 mm
- ^g Requires adapter kit (included)
- $^{\rm h}\,$ System is serviceable while in the rack
- ⁱ Requires additional rail guide (included in kit) for full serviceability
- ^j With middle brackets removed
- ^k With rear brackets removed (applies to 2-post or cantilever mount only)
- ¹ SRB is staged furthest to the rack door
- ^m Rail depth is dependent on spacing between the front and rear mounting flanges of the rack Add amount below based on flange type:
 - Square hole (5.7mm)
 - Round hole (11.8mm)
- ⁿ Rail depth represents cabinet assembly only and does not represent inner rail component that attaches to chassis
- ° Footnote intentionally left blank
- ^p Chassis type utilizes the Self-Adjusting Rail Feature to install properly into rack
- ^q Depth maybe greater based on rail adjustability range
- $^{\rm r}$ Rail threaded mount only compatible with #10-32 thread type
- ^s Rail is only intended for use with ruggedized transit case (Pelican custom rack 25-036329-01)
- ^t Rail supports partial or full in rack service position
- " Requires swap screws (included in Rail Kit), based on chassis rack ear type and Rack Installation guide
- $^{\scriptscriptstyle v}$ The hole diameter of the threaded hole rack flange is required to be greater than 4mm
- $^{\scriptscriptstyle W}$ The hole diameter of the threaded hole rack flange equal or greater than 10-32UNF-2B





Dell Technologies Enterprise Systems Rail Sizing and Rack Compatibility Matrix Table 3. Dell Technologies rack compatibility matrix

						ll-branded APC Racks 100X717/AR3104X717)	Dell xx20	Dell xx20D/xx20S	Dell xx20W	Dell xx10	HP 10XXX	HP/Compaq 9XXX	IBM S2	APC Netshelter SX m Wide x 1070mm Deep)	'Post Rack Spacing TITAN-D/ TITAN S	'Post Rack Spacing TITAN /TITAN SS	iebert Foundation	atsworth Teraframe	ightline Vantage S2
		Product	Rail Identifier	Mounting Interface	Rail Type	De (AR3								(600m	24,	24'		Ċ	٨r
			Α7	ReadyRails II	Sliding	√ ²	~	×	~	×	~	√1	~	√ ²	x	X	~	~	~
		R640 (8-HDD)	A8	ReadyRails	Stab-in Static	×	*	×	*	×	×	×	*	~	✓ ¹⁵	√ ¹⁵	×	~	×
			A10	Generic Tool-less	Drop-in/ Stab-in	✓2	✓	~	~	~	~	✓1	~	√ ²	√14	√ 14	~	~	~
			Α7	ReadyRails II	Sliding	√3,4	√ ²	√ 2	<	√9	~	√1	*	√ ^{3,4}	x	x	~	~	~
		R640 (4-HDD/10-HDD)	A8	ReadyRails	Stab-in Static	~	~	1	✓	×	*	*	~	~	√ ¹⁵	√ ¹⁵	~	~	~
	<		A10	Generic Tool-less	Drop-in/ Stab-in	√3,4	√ ²	√ ²	✓	√9	×	✓1	✓	√ ^{3,4}	√ ¹⁴	✓ ¹⁴	✓	~	✓
ERS	:dge ^{TA}	R6525 (8-HDD)	A15	ReadyRails II	Sliding	√4,12	√ ²	√ ²	✓	×	✓	✓1	✓	√3,4	x	x	✓	~	~
SERV	werE	R650 (8-HDD) R660/R6615/R6625	A14	ReadyRails	Stab-in Static	×	*	~	<	~	~	~	<	~	✓ ¹⁵	√ ¹⁵	×	~	~
	Pc	(0-HDD/8-HDD)	A16	Generic Tool-less	Drop-in/ Stab-in	√4,12	√ 2	√2	~	×	~	√1	~	√3,4	√14	√14	×	~	~
		R6525 (4-HDD/10-HDD)	A15	ReadyRails II	Sliding	√4,12	√4,12	√4,12	~	√ ¹³	~	✓1	√ ¹³	√ ^{3,4}	x	x	×	~	√ ¹³
		R660/ R6615/R6625	A14	ReadyRails	Stab-in Static	×	~	~	~	×	~	✓	~	~	√ ¹⁵	√ ¹⁵	×	~	×
		R670/R6725/R6715/R470	A16	Generic Tool-less	Drop-in/ Stab-in	√4,12	√4,12	√4,12	✓	√ ¹³	~	✓1	√ ¹³	√ ^{3,4}	√14	√ ¹⁴	✓	~	√ ¹³
		R340 (8-HDD)	A12	ReadyRails II	Sliding	√ ^{3,4}	√ ²	√ ²	×	√9	×	✓1	×	√3,4	X	x	✓	~	×
		R360 (8-HDD)	A8	ReadyRails	Stab-in Static	×	✓	✓	✓	✓	~	×	~	~	√ ¹⁵	√ 15	~	~	×

22

D&LLTechnologies



					l-branded APC Racks 100X717/AR3104X717)	Dell xx20	ell xx20D/xx20S	Dell xx20W	Dell xx10	HP 10XXX	IP/Compaq 9XXX	IBM S2	APC Netshelter SX n Wide x 1070mm Deep)	Post Rack Spacing ITAN-D/ TITAN S	Post Rack Spacing TITAN /TITAN SS	ebert Foundation	tsworth Teraframe	ghtline Vantage S2	TEO
_	Product	Rail Identifier	Mounting Interface	Rail Type	Del (AR3						- -		(600mr	24" T	24".	Li	Cha	Wri	
	R340 (4-HDD)	A12	ReadyRails II	Sliding	√ ^{3,4}	√ ²	√ ²	*	√9	<	✓1	~	√3,4	x	x	*	*	~	
	R350 (4-HDD) R360 (4-HDD)	A8	ReadyRails	Stab-in Static	~	~	~	~	*	1	~	~	*	√ ¹⁵	√ 15	*	*	~	
	R440 (8-HDD) R450 (8-HDD) R6415 (8-HDD)	A8	ReadyRails	Stab-in Static	~	1	1	*	~	~	*	*	~	√ ¹⁵	√ 15	~	~	~	
	R6515 (8-HDD) R650xs (0-HDD/8-HDD) R660xs(0-HDD/8-HDD	A11	Generic Tool-less	Drop-in/ Stab-in	√2	*	*	*	<	<	✓1	*	√ ²	√ ¹⁴	√14	*	*	*	
	R440 (4-HDD/10-HDD) R450 (4-HDD) R6415 (4-HDD/10-HDD)	A8	ReadyRails	Stab-in Static	~	~	~	~	*	*	~	*	*	√ ¹⁵	√ ¹⁵	*	*	1	
	R6515 (4-HDD/10-HDD) R650xs (4-HDD/10-HDD) R650xs (8-HDD NVME) R660xs(4-HDD/10-HDD)	A11	Generic Tool-less	Drop-in/ Stab-in	√ 3,4	√2	√2	*	√9	✓	✓1	*	√3,4	√14	√14	*	*	*	
	R540/R540xd	B6	ReadyRails II	Sliding	√ ²	1	1	×	~	×	√1	×	√2	x	X	~	✓	~	
	R740/R740xd/ R7415/R7425/R7515	B4	ReadyRails	Stab-in Static	×	*	×	×	×	*	×	*	~	√ ¹⁵	√ ¹⁵	*	~	×	
		B13	Generic Tool-less	Drop-in/ Stab-in	✓2	~	~	~	✓	✓	✓1	✓	√ ²	√ ¹⁴	√14	~	✓	~	
	R550	B21	ReadyRails II	Sliding	√2	~	~	 	×	✓	✓1	✓	√ ²	x	x	✓	1	×	
	R750xs R760xs	B20	ReadyRails	Stab-in Static	~	~	~	×	✓	✓	×	~	~	√ ¹⁵	√ 15	✓	✓	 Image: A start of the start of	
		B22	Generic Tool-less	Drop-in/ Stab-in	✓2	~	~	×	~	~	✓1	~	√ ²	✓ ¹⁴	✓ ¹⁴	~	✓	 Image: A second s	

23

D&LLTechnologies



				ll-branded APC Racks 100X717/AR3104X717)	Dell xx20	Jell xx20D/xx20S	Dell xx20W	Dell xx10	HP 10XXX	HP/Compaq 9XXX	IBM S2	APC Netshelter SX m Wide x 1070mm Deep)	'Post Rack Spacing TITAN-D/ TITAN S	 Post Rack Spacing TITAN /TITAN SS 	iebert Foundation	atsworth Teraframe	ightline Vantage S2
Product	Rail Identifier	Mounting Interface	Rail Type	De (AR3								(600m	24'.	24'		Chi	٨٢
	B21	ReadyRails II	Sliding	√ 3,5,12	√ ²	√ ²	*	×	×	✓1	×	√ ^{3,4}	x	x	✓	×	×
R760/R7615/R7625 R770/R7715/R7725/R570	B20	ReadyRails	Stab-in Static	~	*	*	*	×	×	×	×	~	✓ ¹⁵	✓15	✓	×	×
	B22	Generic Tool-less	Drop-in/ Stab-in	√ ^{3,5,12}	√ ²	√ ²	*	×	×	✓1	×	√ ^{3,4}	✓ ¹⁴	✓14	✓	×	×
	B6	ReadyRails II	Sliding	√ ^{3,5,12}	√ ²	√ ²	*	×	×	✓1	*	√ ^{3,4}	x	x	~	×	×
R7525 R750	B4	ReadyRails	Stab-in Static	×	~	~	~	×	×	*	×	~	✓15	✓15	1	*	 Image: A set of the set of the
	B13	Generic Tool-less	Drop-in/ Stab-in	√ 3,5,12	√ ²	√ ²	~	×	×	✓1	×	√3,4	✓14	✓14	×	×	~
P750v 2	B19	Generic Tool-less	Drop-in/ Stab-in	√4,6,12	✓ ⁵	✓ ⁵	~	√ 10,13	√ 10,13	√ 10,13	√10	~	√ ¹⁴	√ 14	1	√ 10,13	√ 10,1 3
K750Xa	B17	ReadyRails	Stab-in Static	√4,6	√4	√4	*	√ 10	√ 10	√ 10	√ 10	~	√ 15	√ 15	✓	√ 10,13	✓10,1 3
P760va	B25	Generic Tool-less	Drop-in/ Stab-in	√4,7,12	√4,7	×	~	√ 10,13	√ 10,13	√ ^{10,13}	√ 10,13	√4,7,12	√ ¹⁴	✓13,14	√ ¹³	√ 10,13	✓10,1 3
K7 OOXa	B33	ReadyRails	Stab-in Static	√4,7	√4,7	×	1	√ 10	√ ¹⁰	√ ¹⁰	√ 10	√4,7	√14	√ 14	✓	√ 10	√ ¹⁰
R770 Cold Aisle	B35	ReadyRails	Stab-in, Static	√4,7	√4,7	~	~	√10	√10	√ 10	√10	√4,7	x	X	✓	√10	√10
R670 Cold Aisle	A28	ReadyRails	Stab-in, Static	√4,7	√4,7	~	~	√10	√ 10	√ 10	√ 10	√4,7	x	x	✓	√ 10	√ ¹⁰
R740xd2	-	Generic Tool-less	L-Bracket Static	~	✓	~	✓	*	*	*	~	~	√14	√14	✓	*	✓

24

D&LLTechnologies



				ll-branded APC Racks 100X717/AR3104X717)	Dell xx20	Jell xx20D/xx20S	Dell xx20W	Dell xx10	HP 10XXX	HP/Compaq 9XXX	IBM S2	APC Netshelter SX m Wide x 1070mm Deep)	' Post Rack Spacing TITAN-D/ TITAN S	' Post Rack Spacing TITAN /TITAN SS	iebert Foundation	atsworth Teraframe	ightline Vantage S2
Product	Rail Identifier	Mounting Interface	Rail Type	De (AR3								(600m	24"]	24"		ĊŸ	Wr
R760xd2	B23	Generic Tool-less	Stab-in Sliding	√4,7,12 ,13	√ 4,7,12 13	*	√ ¹³	✓ 10,11 • 13,17	√ 10,13	√ 10,13	√10,13	✓4,7,12 ,13	x	x	√ ¹³	√10,13	√ 10,1 3
R840	B15	Generic Tool-less	Drop-in/ Stab-in	√4,6,12	√5	√5	~	√ ¹⁰	√ ¹⁰	√ ^{10,13}	√ ¹⁰	~	x	x	✓	√10	√ 10,1 3
50/0	B19	Generic Tool-less	Drop-in/ Stab-in	√ 4,7,12 ,13	√7,13	~	√ ¹³	√ 10,13	√ 10,13	√ 10,13	√ 10,13	√4,7,12 ,13	√14	✔ 10,13,14	√ ¹³	√ ¹³	√ 10,1 3
K860	B24	ReadyRails	Stab-in Static	√4,6	√6	~	~	√ 10	√ ¹⁰	√10	√ 10	√4,6	√ ¹⁵	√ 15	1	~	√10
R940	B12	ReadyRails II	Sliding	√3,6,12	√ 3,6,12	√ ^{3,6,12}	~	√ ¹³	✓13	√ ¹³	√ ¹³	~	√ ¹⁵	√ 15	1	~	~
R960	B25	Generic Tool-less	Drop-in/ Stab-in	√4,7,12 ,13	√7,13	~	√ 13	√ 10,13	√ 10,13	√ 10,13	√ 10,13	✓14	✓10,13,1 4	✓14	√ ¹³	√ 13	✓10,1 3
R940xa	B16	Generic Tool-less	Drop-in/ Stab-in	√ ^{4,6,12}	√5	✓ ⁵	~	√ 10	√ 10	√ ^{10,13}	√ 10	~	x	x	1	√ 10	✓10,1 3
EV2/EV2-	B10	ReadyRails II	Sliding	√4,6,12	✓ ⁵	✓ ⁵	~	✓10	√ ¹⁰	√ 10,13	√ 10	√ 4,6,12	x	x	✓	√ 10	✓10,1 3
FAZ/FAZS	B11	ReadyRails II	Stab-in Static	√4,6	~	~	~	√10	√10	√ 10	✓10	√4,6	x	x	1	✓10	√10
C4130/C4140	Α9	ReadyRails II	Stab-in Static	√7	√ 4,7,10	√ 4,7,10	√10	x	x	x	x	√7	X	X	√ 10	X	X
VRTX	C3	ReadyRails II	Sliding	√2	~	~	1	1	1	✓1	~	√ ²	√ 15	√15	1	~	~
R240/R250/R260	A4	ReadyRails	Stab-in Static	~	~	×	×	~	~	~	~	~	√ 15	√ 15	✓	~	✓
M1000e	-	RapidRails	L-Bracket Static	√4,5	~	~	~	~	~	~	×	√ ^{4,5}	x	X	1	~	~
miouce	-	VersaRails	L-Bracket Static	√4,5	✓	✓	✓	✓	✓	✓	✓	√ 4,5	Х	X	1	×	✓

25

D&LLTechnologies



						ll-branded APC Racks 100X717/AR3104X717)	Dell xx20)ell xx20D/xx20S	Dell xx20W	Dell xx10	HP 10XXX	HP/Compaq 9XXX	IBM S2	APC Netshelter SX n Wide x 1070mm Deep)	Post Rack Spacing TITAN-D/ TITAN S	Post Rack Spacing TITAN /TITAN SS	ebert Foundation	itsworth Teraframe	ightline Vantage S2
		Product	Rail Identifier	Mounting Interface	Rail Type	Del (AR3						-		(600mr	24" T	24"	Ľ	Cha	Wri
		MX7000	C5	ReadyRails II	L-Bracket Static	√ 4,6,16	√ 16	✓16	√ 16	√ 10,16	√ 10	√ ¹⁰	√ 10	√ 4,6,16	√ 15	√15	*	√ 10	√ 10
		XR2	A12	ReadyRails II	Sliding	√3,4	√ ²	√ ²	~	√9	~	✓1	✓	√3,4	x	x	×	*	~
			Α4	ReadyRails	Stab-in Static	*	~	~	~	~	~	~	~	~	✓15	√ ¹⁵	*	*	~
			-	Generic	Stab-in Static ^t	-	-	-	-	-	-	-	-	-	-	-	-	-	-
			A20	Generic Tool-less	Stab-in Sliding	~	~	~	✓	~	✓	~	✓	×	√ ¹⁴	√ 14	✓	×	~
		XR11/XR12	A21 ¹⁸	Generic Tool-less	Stab-in Sliding	-	-	-	-	-	-	-	-	-	-	-	-	-	-
		C1100	-	Generic Tool-less	Sliding	~	~	~	✓	~	✓	✓	✓	×	x	x	✓	×	~
		C2100	-	Generic	Sliding	~	~	~	✓	~	✓	~	✓	~	x	X	✓	*	~
		C410x	-	VersaRails	Stab-in Static	√ ⁸	√ 8	√ ⁸	√ ⁸	√ ⁸	~	~	√ ⁸	√ 8	x	x	X	x	~
	PowerEdge C	С5ххх	-	Generic Tool-less	L-Bracket Static	~	~	~	×	~	*	×	×	×	x	X	×	×	✓
		C64xx C65xx C66xx	-	Generic Tool-less	L-Bracket Static	~	~	~	~	√16	~	~	~	*	√14	√14	~	*	~
		C8000	-	Generic Tool-less	L-Bracket Static	√4	✓	~	✓	✓	✓	✓	✓	✓4	X	X	✓	×	×
			-	Generic Tool-less	Sliding	√4,6	√ 4, 11	√ 4, 11	✓11	*	1	*	1	√4,6	X	X	×	×	~

D&LLTechnologies



					ll-branded APC Racks 100X717/AR3104X717)	Dell xx20	Jell xx20D/xx20S	Dell xx20W	Dell xx10	HP 10XXX	HP/Compaq 9XXX	IBM S2	APC Netshelter SX m Wide x 1070mm Deep)	' Post Rack Spacing FITAN-D/ TITAN S	 Post Rack Spacing TITAN /TITAN SS 	ebert Foundation	itsworth Teraframe	ightline Vantage S2
	Product	Rail Identifier	Mounting Interface	Rail Type	Del (AR3						-		(600m)	24" 1	24"	5	Cha	Wri
	HS5620	B21	ReadyRails II	Sliding	√2	*	~	*	~	~	√1	~	√2	x	x	~	~	~
		B20	ReadyRails	Stab-in Static	*	1	~	*	~	*	~	~	~	√ 15	√15	~	~	~
e		B22	Generic Tool-less	Drop-in/ Stab-in	✓2	~	1	×	×	×	✓1	×	√ 2	√ ¹⁴	√14	~	~	~
	HS5610 Cold Aisle	A22	ReadyRails	Stab-in, Static	~	✓	~	~	✓	*	×	✓	~	√ ¹⁵	√ 15	~	~	~
		A8	ReadyRails	Stab-in Static	~	~	~	~	~	~	~	~	~	✓15	√ ¹⁵	~	~	~
	133010(0-1100)	A11	Generic Tool-less	Drop-in/ Stab-in	√ ²	*	×	*	✓	×	✓1	*	√ ²	√14	✓ ¹⁴	×	×	~
	HS5610(4-HDD/10-HDD)	A8	ReadyRails	Stab-in Static	×	*	~	*	×	*	×	*	×	√15	√ 15	~	×	×
		A11	Generic Tool-less	Drop-in/ Stab-in	√3,4	√2	√ ²	*	√9	*	√1	*	√3,4	√14	√14	~	~	~
	XR4000R	A23	Generic Tool-less	Stab-in Sliding	×	~	~	*	*	*	*	*	×	~	~	~	~	~
		A24 ¹⁸			X	X	X	X	X	X	X	X	X	X	X	X	X	X
9		A25 ¹⁸			x	X	x	x	x	x	x	X	X	x	x	x	X	x
	YBSUUD	B31 ¹⁸	Generic Tool-less	L-Bracket Static	X	X	X	X	X	X	X	X	X	X	X	X	X	X
		B36 ¹⁸	Generic Tool-less	L-Bracket Static	X	X	X	X	X	X	X	X	X	X	X	X	X	X
	¥85640	A26	Generic Tool-less	Stab-in Sliding	~	~	~	~	~	~	~	~	~	~	~	~	✓	~
	AKJOIU	A27 ¹⁸			x	x	x	x	X	x	x	X	x	x	x	x	x	X

27

D&LLTechnologies


Dell Technologies Enterprise Systems Rail Sizing and Rack Compatibility Matrix

					ll-branded APC Racks 100X717/AR3104X717)	Dell xx20)ell xx20D/xx20S	Dell xx20W	Dell xx10	HP 10XXX	HP/Compaq 9XXX	IBM S2	APC Netshelter SX n Wide x 1070mm Deep)	Post Rack Spacing TITAN-D/ TITAN S	Post Rack Spacing TITAN /TITAN SS	ebert Foundation	itsworth Teraframe	ightline Vantage S2
	Product	Rail Identifier	Mounting Interface	Rail Type	Del (AR3						- -		(600mr	24" T	24".	E	Cha	Wri
	XR7620	B29	Generic	Stab-in	~	~	~	~	~	~	~	~	✓	*	~	✓	*	~
		B30 ¹⁸	Tool-less	Sliding	X	X	X	X	X	X	X	X	X	X	X	X	X	X
	XE2420	-	Generic Tool-less	Stab-in Static Standard	√17	√17	√17	√17	√17	√17	√17	√17	√17	x	x	√17	√17	√17
	XE8640	B27	Generic Tool-less	Stab-in Sliding	√ 6,12	√7	~	~	√13	√13	√ 13	√13	√6,13	√14	√13,14	*	√ 13	√ ¹³
XE	XE9680	B28	Generic Tool-less	Stab-in Sliding	x	x	√4,12	x	x	x	x	x	x	√ 4,12, 14	x	x	x	x
PowerEdge	XE9640	B25	Generic Tool-less	Drop-in/ Stab-in	√4,7,12 ,13	√4,7	~	~	√10,13	√10,13	√10,13	√10,13	√ 4,7,12 ,13	√14	√13,14	✓	√ 10,13	√10,1 3
	XE9680L/XE9685L	B38	Generic Tool-less	Stab-in Sliding	x	x	√4,7	x	x	x	x	x	x	√4,14	x	x	x	x
	XE7740/XE7745	B37	Generic Tool-less	Stab-in Sliding	x	x	√4,6	x	x	x	x	x	x	√ 4,14	x	x	x	x

28

D&LLTechnologies



Dell Technologies Enterprise Systems Rail Sizing and Rack Compatibility Matrix

						ll-branded APC Racks 100X717/AR3104X717)	Dell xx20)ell xx20D/xx20S	Dell xx20W	Dell xx10	HP 10XXX	HP/Compaq 9XXX	IBM S2	APC Netshelter SX n Wide x 1070mm Deep)	Post Rack Spacing TITAN-D/ TITAN S	Post Rack Spacing TITAN /TITAN SS	ebert Foundation	itsworth Teraframe	ightline Vantage S2
		Product	Rail Identifier	Mounting Interface	Rail Type	Del (AR3						-		(600m)	24" 1	24"	5	Cha	Wri
	1	1081AD/2161AD/1082DS /2162DS 4322DS	А5	ReadyRails	Stab-in Static	~	~	~	~	*	*	*	~	~	-	-	✓	*	~
	KVN	180AS/2160AS 2161DS/2161DS- 2/4161DS 2321DS	-	Generic	Stab-in Static	*	*	~	~	~	~	~	*	*	-	-	~	~	*
VITCHES		PC8132/PC8132F	Α5	ReadyRails	Stab-in Static	~	~	~	~	~	~	~	~	*	-	-	~	~	~
SWI	PC8164	PC8164/PC8164F																	
	Networki	S4820T/S6000	А5	ReadyRails	Stab-in Static	~	~	~	*	*	*	*	~	~	-	-	*	*	~
		\$5000	-	Generic	Stab-in Static	*	1	*	~	~	*	<	*	*	-	-	<	<	~
		R7910	B6	ReadyRails II	Sliding	√2	1	×	×	~	✓	√1	×	√2	-	-	~	×	~
			B4	ReadyRails	Stab-in Static	×	1	1	✓	✓	×	~	~	~	-	-	~	×	×
	SNOL	Precision 3930 Rack	Α4	ReadyRails	Stab-in Static	~	×	~	×	×	×	×	×	×	-	-	*	×	~
	WORKSTATIO		B6	ReadyRails II	Sliding	√ ²	✓	~	×	✓	×	✓1	×	√ ²	x	x	✓	×	~
			B4	ReadyRails	Stab-in Static	✓	✓	✓	×	×	×	×	×	~	√ ¹⁵	√ 15	√	×	 Image: A start of the start of
		Precision 7920 Rack	B13	Generic Tool-less	Drop-in/ Stab-in	√2	~	~	~	*	*	√1	~	√2	√14	√14	*	*	~

29

D&LLTechnologies



					1-branded APC Racks 100X717/AR3104X717)	Dell xx20	ell xx20D/xx20S	Dell xx20W	Dell xx10	HP 10XXX	IP/Compaq 9XXX	IBM S2	APC Netshelter SX n Wide x 1070mm Deep)	Post Rack Spacing IITAN-D/ TITAN S	Post Rack Spacing TITAN /TITAN SS	ebert Foundation	tsworth Teraframe	ghtline Vantage S2	(E)
	Product	Rail Identifier	Mounting Interface	Rail Type	Del (AR3						-		, (600mr	24" T	24"	Li	Cha	Wri	
		B21	ReadyRails II	Sliding	√ ^{3,5,12}	√ ²	√ ²	×	×	*	✓1	×	√3,4	x	x	~	~	×	
	Precision 7960 Rack	B20	ReadyRails	Stab-in Static	~	~	*	~	×	*	×	×	×	✓15	√15	~	~	×	
		B22	Generic Tool-less	Drop-in/ Stab-in	√ 3,5,12	√ ²	√ ²	~	~	~	✓1	*	√3,4	√14	√14	*	✓	 Image: A set of the set of the	
	T7600/T7610	C2	ReadyRails II	Sliding	✓2	✓11	✓11	✓11	~	×	✓1	*	✓2	-	-	*	*	~	ĺ
	R5500/R7610	B2	ReadyRails	Sliding	√3	√ ²	√ ²	~	~	×	✓1	×	✓3	-	-	~	✓	×	
	FPM185 (without KVM)		ReadyRails II	Sliding	~	~	~	~	~	~	~	×	~	-	-	x	*	~	
KMM	FPM185 (with KVM)	-	ReadyRails II	Sliding	~	~	1	~	~	1	1	~	✓	-	-	x	~	~	
	1750	-	RapidRails	Sliding	1	~	×	~	~	×	× -	×	× -	-	-	1	1	✓	
	1766	-	VersaRails	Sliding	~	1	*	~	~	*	*	*	×	-	-	*	*	×	
UPS	Dell Rack Mount UPS Family	В5	ReadyRails	Stab-in Static	~	~	*	~	~	*	~	*	~	-	-	~	~	~	
OTHER	1U Fixed Equipment Shelf	Α4	ReadyRails	Stab-in Static	~	~	~	~	~	*	*	*	*	-	-	~	~	~	
STORAGE	PROXPE1 PROXPE2	B14	-	L-Bracket Static	√ 6,13,1 6	~	*	✓13	✓13	√13	√13	√13	~	~	*	√ 13	√13	√ 13	

D&LLTechnologies



					ll-branded APC Racks 100X717/AR3104X717)	Dell xx20	Jell xx20D/xx20S	Dell xx20W	Dell xx10	HP 10XXX	HP/Compaq 9XXX	IBM S2	APC Netshelter SX m Wide x 1070mm Deep)	'Post Rack Spacing TITAN-D/ TITAN S	'Post Rack Spacing TITAN /TITAN SS	iebert Foundation	atsworth Teraframe	ightline Vantage S2
	Product	Rail Identifier	Mounting Interface	Rail Type	De (AR3								(600m	24'.	24'		ĊŸ	Wr
	NX3300/NX400	Α7	ReadyRails II	Sliding	√ ²	*	*	×	×	×	✓	*	X	x	x	~	~	✓
	NX3300/NX400	A8	ReadyRails	Stab-in Static	×	*	×	✓	✓	✓	*	*	√ ¹⁵	√ ¹⁵	√15	✓	~	✓
	NX3200	B6	ReadyRails II	Sliding	✓2	✓	×	×	×	×	✓1	×	x	x	x	✓	~	✓
	142200	B4	ReadyRails	Stab-in Static	~	*	*	✓	✓	✓	*	*	√ ¹⁵	√ ¹⁵	√ ¹⁵	✓	~	~
		A3	ReadyRails	Sliding	√ ²	✓	~	×	*	*	✓1	×	x	x	x	*	~	~
	NX3500 Controller	A4	ReadyRails	Stab-in Static	~	✓	×	✓	✓	✓	×	×	√ ¹⁵	√ ¹⁵	√ ¹⁵	✓	×	~
א מחור	NX3500 UPS	Α4	ReadyRails	Stab-in Static	~	✓	✓	✓	✓	✓	×	×	✓	√ ¹⁵	√ 15	~	×	~
0.44	DY(000C	A4	ReadyRails	Stab-in Static	~	✓	~	~	*	*	~	~	~	√ 15	√ 15	~	~	~
	DX6000G	A6	ReadyRails	Stab-in Static	~	~	~	1	×	×	~	~	*	x	x	~	~	~
	NY200/DY60045	A3	ReadyRails	Sliding	√ ²	×	×	~	×	×	✓1	*	√2	x	x	✓	~	✓
	147200/0700042	A4	ReadyRails	Stab-in Static	~	✓	~	✓	✓	✓	✓	~	✓	√ ¹⁵	√ ¹⁵	✓	~	~
	NX3000/DX6000	B1	ReadyRails	Sliding	√ ²	✓	~	✓	✓	✓	✓1	✓	√ ²	X	X	~	~	~
	11,3000,270000	A2	ReadyRails	Stab-in Static	~	✓	~	×	×	×	×	~	×	✓15	✓15	✓	~	×
	NX3100/DL2200/	В3	ReadyRails	Sliding	✓2	~	*	~	~	~	✓1	*	✓2	x	x	~	~	~

D&LLTechnologies



				ll-branded APC Racks 100X717/AR3104X717)	Dell xx20	oell xx20D/xx20S	Dell xx20W	Dell xx10	HP 10XXX	HP/Compaq 9XXX	IBM S2	APC Netshelter SX n Wide x 1070mm Deep)	Post Rack Spacing TITAN-D/ TITAN S	Post Rack Spacing TITAN /TITAN SS	ebert Foundation	itsworth Teraframe	ightline Vantage S2	TET
Product	Rail Identifier	Mounting Interface	Rail Type	Del (AR3						Ţ		(600mr	24" T	24"	Ċ	Cha	Wri	
DX6012S/DR4000	B4	ReadyRails	Stab-in Static	~	×	×	×	×	×	×	×	×	√ 15	√ 15	✓	~	×	
MD3060e/MD3660	-	VersaRails	L-Bracket Static	√4,6	✓4	✓4	×	✓10	√ 10	x	√ 10	√4,6	x	x	✓	√10	х	
MD12xx/14xx/32xx/36xx	В9	ReadyRails II	L-Bracket Static	~	✓	×	×	1	×	*	×	1	√ 15	√ ¹⁵	1	~	1	
NV26vv	-	RapidRails	L-Bracket Static	~	✓	×	×	✓	×	x	✓	×	x	x	x	~	✓	
NX 30XX	-	VersaRails	L-Bracket Static	~	×	×	×	~	✓	×	×	~	X	X	✓	✓	×	
MD1120	-	RapidRails	L-Bracket Static	~	~	~	~	~	~	X	~	~	X	X	X	✓	~	
MUTIZU	-	VersaRails	L-Bracket Static	~	~	~	~	~	~	~	~	~	X	X	✓	~	×	
MD1000/MD3000	-	RapidRails	L-Bracket Static	~	~	✓	~	~	✓	✓	✓	~	X	X	✓	✓	~	
WD 1000/WD3000	-	VersaRails	L-Bracket Static	~	~	~	~	~	~	✓	~	~	X	X	✓	✓	~	
	Β7	ReadyRails	Stab-in Static	~	~	~	×	×	~	~	1	×	√ 15	√ ¹⁵	~	~	~	
PV114T/PV114X	-	RapidRails	Sliding	✓2	✓	~	~	~	~	✓1	~	✓2	X	x	~	~	1	
	-	VersaRails	Sliding	√ ²	✓	×	×	✓	×	✓1	~	√ 2	X	x	~	~	1	
	-	RapidRails	L-Bracket Static	~	1	~	~	~	~	~	×	~	X	x	~	~	~	
PV124T	-	VersaRails	L-Bracket Static	~	~	~	~	*	*	*	~	*	x	x	*	*	~	

D&LLTechnologies



Dell Technologies Enterprise Systems Rail Sizing and Rack Compatibility Matrix

					ll-branded APC Racks 100X717/AR3104X717)	Dell xx20	oell xx20D/xx20S	Dell xx20W	Dell xx10	HP 10XXX	HP/Compaq 9XXX	IBM S2	APC Netshelter SX n Wide x 1070mm Deep)	Post Rack Spacing TITAN-D/ TITAN S	Post Rack Spacing TITAN /TITAN SS	ebert Foundation	ıtsworth Teraframe	ightline Vantage S2
	Product	Rail Identifier	Mounting Interface	Rail Type	Del (AR3						- -		(600mr	24" T	24".	Li	Cha	Wri
	ES7500 Controllor	A1	ReadyRails	Sliding	✓3	✓2	√ ²	*	~	*	√1	*	√3	x	x	*	~	~
	rs7500 Controller	A2	ReadyRails	Stab-in Static	~	~	~	*	*	*	*	*	~	√15	√15	*	*	~
	FS7500 UPS	A4	ReadyRails	Stab-in Static	~	~	~	*	*	*	*	*	~	√ 15	✓15	*	~	~
		В9	ReadyRails II	L-Bracket Static	~	~	~	~	×	×	×	✓	✓	√ ¹⁵	√ 15	✓	×	~
ıalLogic	FS76xx/PS41xx/PS61xx	-	RapidRails	L-Bracket Static	~	~	~	~	~	*	x	✓	✓	x	x	x	~	~
Equ		-	VersaRails	L-Bracket Static	~	~	~	~	~	~	~	~	✓	x	x	*	~	~
	PS6500/6510	-	ReadyRails	Sliding	✓7	✓2	√2	~	*	*	x	*	√7	√15	√ 15	*	*	*
	PS4000/6000/6010	-	Generic	L-Bracket Static	*	*	~	~	~	~	~	~	~	x	x	*	*	*
Dell Compellent	SC20xx/SC40xx	-	Generic	L-Bracket Static	*	*	~	*	*	*	*	*	*	√15	√15	*	*	*

33

D&LLTechnologies



	-			ll-branded APC Racks 100X717/AR3104X717)	Dell xx20	Dell xx20D/xx20S	Dell xx20W	Dell xx10	HP 10XXX	HP/Compaq 9XXX	IBM S2	APC Netshelter SX m Wide x 1070mm Deep)	'Post Rack Spacing TITAN-D/ TITAN S	'Post Rack Spacing TITAN /TITAN SS	iebert Foundation	atsworth Teraframe	ightline Vantage S2
Product	Rail Identifier	Mounting Interface	Rail Type	De (AR3								m00ð)	24'.	24'		ĊŸ	Wr
	B6	ReadyRails II	Sliding	√ ²	✓	✓	✓	×	×	√1	~	√ ²	√ ¹⁵	√ 15	✓	~	~
SC8000	Β4	ReadyRails	Stab-in Static	*	1	*	*	*	*	*	1	*	√15	√15	✓	~	*
	В9	ReadyRails II	L-Bracket Static	~	~	~	×	×	×	×	*	×	x	x	*	~	✓
SC2xx/FS86xx	-	RapidRails	L-Bracket Static	×	×	×	×	*	*	X	*	*	X	x	X	~	✓
	-	VersaRails	L-Bracket Static	×	~	×	~	×	~	~	~	~	x	x	~	~	~
SCV30xx/SC50xx/SC7020	B9	ReadyRails II	L-Bracket Static	×	~	×	×	×	*	*	~	*	√ ¹⁵	✓15	~	~	✓
Series 40	-	Generic	Sliding	~	~	~	*	~	~	*	~	~	*	~	~	~	~
Fibre Channel	-	Generic	Stab-in Static	×	~	×	×	*	*	*	*	*	*	~	~	~	✓
SAS (new rails)	-	Generic	Stab-in Static	~	~	*	~	*	*	*	1	*	X	x	~	~	~
SAS (old rails)	-	Generic	Stab-in Static	~	✓	~	✓	×	×	×	~	×	X	X	✓	~	✓
NAS Gen3	-	Generic	Sliding	√ ⁶	×	~	×	×	×	×	×	√ ⁶	x	x	*	~	~

D&LLTechnologies



Notes:

- ¹ A rear door extension kit is required to accommodate the CMA.
- ² CMA may impede access to forward bank of rear-mount PDUs.
- ³ CMA and outer CMA brackets must be removed in order to access the forward bank of rear-mount PDUs.
- ⁴ Rear-mount PDUs may impede extraction of some rear system modules.
- ⁵ The strain relief bar interferes with the forward bank of rear-mount PDUs.
- ⁶ Rails/system block the forward bank of rear-mount PDUs.
- ⁷ Rails/system block both the forward and rearward banks of rear-mount PDUs.
- ⁸ The rear mounting flanges of the rack must be moved rearward.
- ⁹ The CMA tray interferes with rear door lock rod in top U and bottom U.
- ¹⁰ Space for external cable routing is limited.
- ¹¹ May need to adjust the rack's mounting posts back to allow the front door to close.
- ¹² CMA/SRB fully blocks front bank of rear-mount PDUs, and partially blocks the rearward PDU banks. Recommend rotating PDUs 90°.
- ¹³ CMA/SRB must be removed to enable rear door to close for some or all racks in this column
- ¹⁴ The rails align with bezels on Storage systems (unthreaded round-hole rack).
- ¹⁵ The rails require tooled installation for bezel alignment with Storage systems (unthreaded round-hole rack).
- ¹⁶ Strain relief bar might block a portion of the rearward bank of the rear-mount PDUs.
- ¹⁷ Normal Inner rail member allows for tool-less bezel installation and does not enable front rack door to close.
- ¹⁸ Rail is only intended for use within ruggedized transit case (Pelican custom rack 25-036329-01)
- ¹⁹ Rail is only intended for telco short depth rack.

D&LLTechnologies



Integrated Dell Remote Access Controller 9 User's Guide

December 2020 Rev. A02





Notes, cautions, and warnings

(i) NOTE: A NOTE indicates important information that helps you make better use of your product.

CAUTION: A CAUTION indicates either potential damage to hardware or loss of data and tells you how to avoid the problem.

MARNING: A WARNING indicates a potential for property damage, personal injury, or death.

© 2020 Dell Inc. or its subsidiaries. All rights reserved. Dell, EMC, and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be trademarks of their respective owners.

Chapter 1: Overview of iDRAC	16
Benefits of using iDRAC	16
Key features	17
New features added	19
Firmware version 4.40.00.00	19
Firmware version 4.30.30.30	20
Firmware version 4.20.20.20.	21
Firmware version 4.10.10.10.	21
Firmware version 4.00.00.00	22
How to use this guide	23
Supported web browsers	. 23
Supported OS and Hypervisors	23
iDRAC licenses	23
Types of licenses	.23
Methods for acquiring licenses	.24
Acquiring license key from Dell Digital Locker	24
License operations	. 25
Licensed features in iDRAC9	26
Interfaces and protocols to access iDRAC	31
iDRAC port information	34
Other documents you may need	. 35
Contacting Dell	. 35
Accessing documents from Dell support site	36
Accessing Redfish API Guide	. 36
Chapter 2: Logging in to iDRAC	37
Earce Change of Password (ECP)	38
Logging into iDRAC using OpenID Connect	
Logging into iDRAC as local user. Active Directory user, or LDAP user	
Logging in to iDRAC as a local user using a smart card	
Logging in to iDRAC as an Active Directory user using a smart card	ـــــــــــــــــــــــــــــــــــــ
Logging in to iDRAC using Single Sign-On	۵0 III
Logging in to iDRAC SSO using iDRAC web interface	. 40 20
Logging in to iDRAC SSO using CMC web interface	.90 20
Accessing iDRAC using remote RACADM	10 ⊿1
Validating CA certificate to use remote RACADM on Linux	1
	 1
Accessing iDRAC using firmware RACADM	¬⊺ ⊿1
Simple 2-Factor Authentication (Simple 2FA)	¬ . ⊿1
RSA SecuriD 2FA	۰ ۵2
Viewing system health	43
Logging in to iDRAC using public key authentication	43
Multiple iDRAC sessions	0

Contents 3



Resetting default iDRAC password locally	
Resetting default iDRAC password remotely	
Changing the default login password	
Changing the default login password using web interface	
Changing the default login password using RACADM	
Changing the default login password using iDRAC settings utility	47
Enabling or disabling default password warning message	47
Password Strength Policy	
IP Blocking	
Enabling or disabling OS to iDRAC Pass-through using web interface	
Enabling or disabling alerts using RACADM	
Chapter 3: Setting up managed system	
Setting up iDRAC IP address	
Setting up iDRAC IP using iDRAC settings utility	
Setting up iDRAC IP using the CMC web interface	
Auto-discovery	
Configuring servers and server components using Auto Config	
Using hash passwords for improved security	61
Modifying local administrator account settings	
Setting up managed system location	
Setting up managed system location using web interface	
Setting up managed system location using RACADM	
Setting up managed system location using iDRAC settings utility	
Optimizing system performance and power consumption	
Modifying thermal settings using iDRAC web interface	
Modifying thermal settings using RACADM	
Modifying thermal settings using iDRAC settings utility	69
Modifying PCIe airflow settings using iDRAC web interface	
Setting up management station	70
Accessing iDRAC remotely	70
Configuring supported web browsers	
Configuring Internet Explorer	
Configuring Mozilla Firefox	
Configuring web browsers to use virtual console	72
Viewing localized versions of web interface	
Updating device firmware	
Updating firmware using iDRAC web interface	
Scheduling automatic firmware updates	
Updating device firmware using RACADM	
Updating firmware using CMC web interface	
Updating firmware using DUP	
Updating firmware using remote RACADM	
Updating firmware using Lifecycle Controller Remote Services	
Updating CMC firmware from iDRAC	
Viewing and managing staged updates	
Viewing and managing staged updates using iDRAC web interface	
Viewing and managing staged updates using RACADM	
Rolling back device firmware	
Rollback firmware using iDRAC web interface	

	INTE
Rollback firmware using CMC web interface	
Rollback firmware using RACADM	
Rollback firmware using Lifecycle Controller	
Rollback firmware using Lifecycle Controller-Remote Services	
Recovering iDRAC	
Monitoring iDRAC using other Systems Management tools	
Support Server Configuration Profile — Import and Export	
Importing server configuration profile using iDRAC web interface	
Exporting server configuration profile using iDRAC web interface	
Secure Boot Configuration from BIOS Settings or F2	
BIOS recovery	
apter 4: Configuring iDRAC	
Viewing iDRAC information	
Viewing iDRAC information using web interface	
Viewing iDRAC information using RACADM	
Modifying network settings	
Modifying network settings using web interface	
Modifying network settings using local RACADM	
Configuring IP filtering	
Cipher suite selection	
Configuring cipher suite selection using iDRAC web interface	
Configuring cipher suite selection using RACADM	
FIPS mode	
Enabling FIPS Mode	
Disabling FIPS mode	
Configuring services	
Configuring services using web interface	
Configuring services using RACADM	
Enabling or disabling HTTPS redirection	
SEKM Functionalities	98
Using VNC client to manage remote server	99
Configuring VNC server using iDRAC web interface	QQ
Configuring VNC server using RACADM	QQ
Setting un VNC viewer with SSL encryption	100
Setting up VNC viewer without SSL encryption	100
Configuring front nanel display	100
Configuring LCD setting	100
Configuring LOD setting	
Configuring system of LED setting	101 100
Configuring time zone and NTP using iDBAC web interface	102
	IUZ 100
Configuring Line Zone and NTF USING RACADIVI	102 100
Setting first boot device using web interface	IU2
Setting first boot device using web Interface	
Setting first boot device using KACADIVI	
Setting first boot device using virtual console	
Enabling last crash screen	
nabling or disabling OS to iDRAC Pass-through	
Supported cards for US to IDRAC Pass-through	

OTOCO

Enabling or disabling OS to iDRAC Pass-through using web interface	105
Enabling or disabling OS to iDRAC Pass-through using RACADM.	
Enabling or disabling OS to iDRAC Pass-through using iDRAC settings utility.	
Obtaining certificates	107
SSL server certificates	107
Generating a new certificate signing request	108
Automatic Certificate Enrollment	100
	100 100
Viewing server certificate	110
	110
Dowoloading custom SSL certificate signing certificate	110
Doloting custom SSL certificate signing certificate	
Configuring multiple iDBACe using DACADM	
Disabling access to modify iDBAC configuration acttings on heat system	
Disabiling access to modify IDRAC configuration settings on host system	
apter 5: Delegated Authorization using OAuth 2.0	113
apter 6: Viewing iDRAC and managed system information	114
Viewing managed system health and properties	114
Configuring Asset Tracking	114
Viewing system inventory	115
Viewing sensor information	116
Monitoring performance index of CPU, memory, and input output modules	117
Monitoring performance index of CPU, memory, and input output modules using web interfa	ce118
Monitoring performance index for of CPU, memory, and input output modules using RACADI	M 118
Idle Server Detection	118
GPU (Accelerators) Management	
Checking the system for Fresh Air compliance	
Viewing historical temperature data	
Viewing historical temperature data using iDRAC web interface	121
Viewing historical temperature data using RACADM	121
Configuring warning threshold for inlet temperature	122
Viewing network interfaces available on bost OS	100
Viewing network interfaces available on host OS using web interface	100
Viewing network interfaces available on bost OS using RACADM	103
Viewing FlexAddross mozzaning card fabric connections	123
Viewing or terminating iDBAC assessors	107
Terminating iDRAC sessions using web interface	
anter 7: Setting up iDRAC communication	125
Communicating with iDRAC through serial connection using DR9 cable	126
communicating with Drate through solid connection daily DD2 cable	120 126
Configuring BIOS for serial connection	021 126
Configuring BIOS for serial connection	
Configuring BIOS for serial connection Enabling RAC serial connection Enabling IPMI serial connection basic and terminal modes	107
Configuring BIOS for serial connection Enabling RAC serial connection Enabling IPMI serial connection basic and terminal modes	
Configuring BIOS for serial connection Enabling RAC serial connection Enabling IPMI serial connection basic and terminal modes Switching between RAC serial and serial console while using DB9 cable	
Configuring BIOS for serial connection Enabling RAC serial connection Enabling IPMI serial connection basic and terminal modes Switching between RAC serial and serial console while using DB9 cable Switching from serial console to RAC serial	
Configuring BIOS for serial connection Enabling RAC serial connection Enabling IPMI serial connection basic and terminal modes Switching between RAC serial and serial console while using DB9 cable Switching from serial console to RAC serial Switching from RAC serial to serial console	

OTOCO

	INT
Configuring iDRAC to use SOL	
Enabling supported protocol	
Communicating with iDRAC using IPMI over LAN	
Configuring IPMI over LAN using web interface	
Configuring IPMI over LAN using iDRAC settings utility	
Configuring IPMI over LAN using RACADM	
Enabling or disabling remote RACADM	
Enabling or disabling remote RACADM using web interface	
Enabling or disabling remote RACADM using RACADM	
Disabling local RACADM	
Enabling IPMI on managed system	135
Configuring Linux for serial console during boot in RHEL 6	136
Enabling login to the virtual console after boot	
Configuring serial terminal in RHEL 7	
Controlling GRUB from serial console	
Supported SSH cryptography schemes	139
Using public key authentication for SSH	140
Chapter 8: Configuring user accounts and privileges	143
iDRAC user roles and privileges	143
Recommended characters in user names and passwords	144
Configuring local users	
Configuring local users using iDRAC web interface	
Configuring local users using RACADM	145
Configuring Active Directory users	147
Prerequisites for using Active Directory authentication for iDRAC	147
Supported Active Directory authentication mechanisms	148
Standard schema Active Directory overview	148
Configuring Standard schema Active Directory	
Extended schema Active Directory overview	151
Configuring Extended schema Active Directory	
lesting Active Directory settings	
Configuring generic LDAP users	
Configuring generic LDAP directory service using iDRAC web-based interface	
Configuring generic LDAP directory service using RACADM	
Testing LDAP directory service settings	162
Chanter 0. Sustem Configuration Lookdown mode	164
Chapter 9: System Configuration Lockdown mode	
Chanter 40. Configuring DDAC for Single Single On or execut conduction	466
Chapter 10: Configuring IDRAC for Single Sign-On or smart card login	
Prerequisites for Active Directory Single Sign-On or smart card login	
Registering IDRAC on Domain name System	
Creating Active Directory objects and providing privileges	
Creating a Llear in Active Directory for SSO	10/ 167
Generating Kerberos keytab filo	
Configuring iDRAC SSO login for Active Directory users using web interface	100
Configuring iDRAC SSO login for Active Directory users using RACADM	168
Configuring intro CCC login for Active Directory data daling NACADIVI	

Fls. <u>393</u> Mov. <u>32</u>



Enabling or disabling smart card login	. 169
Enabling or disabling smart card login using web interface	169
Enabling or disabling smart card login using RACADM	169
Enabling or disabling smart card login using iDRAC settings utility	169
Configuring Smart Card Login	170
Configuring iDRAC smart card login for Active Directory users	. 170
Configuring iDRAC smart card login for local users	. 170
Using Smart Card to Login	171

Chapter 11: Configuring iDRAC to send alerts	.1/2
Enabling or disabling alerts	172
Enabling or disabling alerts using web interface	172
Enabling or disabling alerts using RACADM	173
Enabling or disabling alerts using iDRAC settings utility	173
Filtering alerts	173
Filtering alerts using iDRAC web interface	173
Filtering alerts using RACADM	174
Setting event alerts	174
Setting event alerts using web interface	174
Setting event alerts using RACADM	174
Setting alert recurrence event	174
Setting alert recurrence events using RACADM	174
Setting alert recurrence events using iDRAC web interface	175
Setting event actions	175
Setting event actions using web interface	175
Setting event actions using RACADM	175
Configuring email alert, SNMP trap, or IPMI trap settings	175
Configuring IP alert destinations	175
Configuring email alert settings	177
Configuring WS Eventing	179
Configuring Redfish Eventing	179
Monitoring chassis events	180
Monitoring chassis events using the iDRAC web interface	180
Monitoring chassis events using RACADM	180
Alerts message IDs	180

Chapter 12: iDRAC 9 Group Manager	
Group Manager	
Summary View	
Network Configuration requirements	
Manage Logins	
Add a New User	
Change User Password	
Delete User	
Configure Alerts	
Export	
Discovered Servers View	
Jobs View	
Jobs Export	

Fis. <u>395</u> Mov. <u>32</u>
 190 00000000000000000000000000000000000
190

Group Information Panel	
Group Settings	
Actions on a selected Server	191
iDRAC Group Firmware Update	

Chapter 13: Managing logs	
Viewing System Event Log	
Viewing System Event Log using web interface	
Viewing System Event Log using RACADM	
Viewing System Event Log using iDRAC settings utility	
Viewing Lifecycle log	
Viewing Lifecycle log using web interface	
Viewing Lifecycle log using RACADM	
Exporting Lifecycle Controller logs	
Exporting Lifecycle Controller logs using web interface	
Exporting Lifecycle Controller logs using RACADM	
Adding work notes	
Configuring remote system logging	
Configuring remote system logging using web interface	
Configuring remote system logging using RACADM	

Chapter 14: Monitoring and managing power in iDRAC	197
Monitoring power	
Monitoring performance index of CPU, memory, and input output modules using web into	erface 197
Monitoring performance index for of CPU, memory, and input output modules using RAC	ADM 198
Setting warning threshold for power consumption	198
Setting warning threshold for power consumption using web interface	
Executing power control operations	
Executing power control operations using web interface	
Executing power control operations using RACADM	
Power capping	199
Power capping in Blade servers	
Viewing and configuring power cap policy	
Configuring power supply options	
Configuring power supply options using web interface	
Configuring power supply options using RACADM	201
Configuring power supply options using iDRAC settings utility	
Enabling or disabling power button	
Multi-Vector Cooling	
Chapter 15: iDRAC Direct Updates	203
Chapter 16: Inventorying, monitoring, and configuring network devices	
Inventorying and monitoring network devices	204
Monitoring network devices using web interface	204
Monitoring network devices using RACADM	204
Connection View	205
Inventorying and monitoring FC HBA devices	

Monitoring FC HBA devices using RACADM. 207 Inventoring SFP Transceiver devices using web interface. 208 Monitoring SFP Transceiver devices using RACADM. 208 Serial Data Capture. 208 Dynamic configuration of virtual addresses. Initiator, and storage target settings. 210 Supported cards for IO Identity Optimization. 212 Virtual or Remote assigned Address and Parsistance Policy behavior when IDRAC is set to Remote Assigned Address and Parsistance Policy behavior when IDRAC is set to Remote Assigned Address and Parsistance Policy behavior when IDRAC is set to Remote Assigned Address and IO Identity Optimization. 212 System behavior for FiexAddress and IO Identity. 213 Enabling or disabiling IO Identity Optimization. 214 SDD Ware Threshold. 214 Configuring persistence policy settings. 215 Chapter 17: Managing storage devices 220 What is RAD. 221 Organizing data storage for availability and performance. 222 Comosing RAD levels 223 Supported enclosures. 236 Supported enclosures using web interface. 236 Supported enclosures (avece suing Wab Interface. 236 Moni		
Montoring FC H1A devices using VACADM. 207 Inventorying and monitoring SFP Transceiver devices 207 Monitoring SFP Transceiver devices using web interface. 208 Monitoring SFP Transceiver devices using RACADM. 208 Stellal Data Capture. 209 Dynamic configuration of virtual addresses, initiator, and storage target settings. 210 Supported cards for IO Identity Optimization. 210 Supported Cards for IO Identity Optimization. 212 Virtual or Remote assigned Address and Persistence Policy behavior when IDRAC is set to Remote Assigned Address and Persistence Policy behavior when IDRAC is set to Remote Assigned Address and IO Identity. 213 Enabling or disabiling IO Identity Optimization. 214 Configuring persistence policy settings. 215 Chapter 17: Managing storage devices 229 Understanding RAD concepts. 220 Organizing Atla storage for availability and performance. 222 Consparing RAD level performance. 228 Supported controllers 230 Supported controllers 230 Supported enclosures. 230 Inderstanding RAD evelops using web interface. </th <th></th> <th>Tree M</th>		Tree M
Inventorying and monitoring SFP Transceiver devices using web Interface	Monitoring FC HBA devices using RACADM	
Monitoring SFP Transcever devices using Web Interface. 208 Monitoring SFP Transcever devices using RACADM 208 Serial Data Capture. 209 Dynamic configuration of Virtual addresses, initiator, and storage target settings. 210 Supported cards for IO Identity Optimization. 211 Virtual or Remote assigned Address and Persistence Policy behavior when IDRAC is set to 212 Virtual or Remote assigned Address and Persistence Policy behavior when IDRAC is set to 213 Enabling or disabiling IO Identity Optimization. 214 SSD Wear Threshold. 214 SD Wear Threshold. 214 Configuring persistence policy settings. 215 Chapter 17: Managing storage devices 220 What is RAID. 221 Organizing data storage for availability and performance. 222 Comparing RAID levels 222 Comparing RAID levels 233 Supported controllers 239 Supported controllers as using with interface. 236 Monitoring storage devices using RACADM. 236 Monitoring storage devices using RACADM. 236 Monitoring storage devices using RACADM. 236 <td< td=""><td>Inventorying and monitoring SFP Transceiver devices</td><td>207</td></td<>	Inventorying and monitoring SFP Transceiver devices	207
Monitoring SH-P Iransceiver devices using RACADM 208 Serial Data Capture. 209 Dynamic configuration of virtual addresses, initiator, and storage target settings. 210 Supported cards for 10 identity Optimization. 210 Supported NIC firmware versions for 10 identity Optimization. 212 Virtual or Remote assigned Address and Persistence Policy behavior when IDRAC is set to Remote-Assigned Address and IO identity. 213 System behavior for FlexAddress and IO identity. 214 SD Wear Threshold. 214 Configuring persistence policy settings. 215 Chapter 17: Managing storage devices 219 Understanding RAID concepts. 221 Organizing data storage for availability and performance. 222 Choosing RAID level performance. 228 Supported enclosures. 230 Supported controllers. 230 Inventorying and monitoring storage devices. 230 Inventorying and monitoring storage devices. 230 Inventorying and monitoring storage devices using RACADM. 236 Monitoring storage devices using RACDM. 236 Monitoring storage devices using RACDM. 236 Monitoring backplane usin	Monitoring SFP Transceiver devices using web interface	
Telemetry Streaming. 209 Byrial Data Copture. 209 Dynamic configuration of virtual addresses, initiator, and storage target settings. 210 Supported cards for IO Identity Optimization. 212 Virtual or Remote assigned Address and Persistence Policy behavior when IDRAC is set to Remote-Assigned Address and Persistence Policy behavior when IDRAC is set to 212 System behavior for FlexAddress and IO Identity. 213 Enabling or disabiling IO Identity Optimization. 214 SDB Wear Threshold. 214 Configuring persistence policy settings. 215 Understanding RAID concepts. 220 What is RAID. 221 Organizing data storage for availability and performance. 222 Comparing RAID level performance. 228 Supported controllers. 229 What is RAID. 221 Organizing data storage for storage devices. 230 Supported controllers. 228 Supported controllers. 228 Supported controllers. 230 Inventorying and montoring storage devices using web interface. 236 Monitoring backplane using IDRAC settings utility. 236 Monit	Monitoring SFP Transceiver devices using RACADM	
Seriel Data Capture 209 Dynamic configuration of virtual addresses, initiator, and storage target settings. 210 Supported eards for IO Identity Optimization. 211 Supported NC firmware versions for IO Identity Optimization 212 Virtual or Remote assigned Address and Persistence Policy behavior when IDRAC is set to Remote-Assigned Address and IO Identity. 213 Enabling or disabling IO Identity Optimization. 214 SD Wear Threshold. 214 Configuring persistence policy settings. 215 Chapter 17: Managing storage devices 219 Understanding RAID concepts. 220 What is RAID. 221 Organizing data storage for availability and performance. 222 Choosing RAID level performance. 228 Supported enclosures. 230 Supported enclosures. 230 Nonitoring storage devices using RACADM. 236 Monitoring storage devices using RACADM. 236 Monitoring storage devices using RACADM. 236 Monitoring storage devices using RACADM. 237 Assigning or unassigning physical disk as global hot spare. 237 Assigning or unassigning physical disk as global hot spare. 237 Assigning or unassigning hysical disk as global hot spare. 237 Assigning or unassigning hysical d	Telemetry Streaming	208
Dynamic configuration of virtual addresses, initiator, and storage target settings. 210 Supported Cards for 10 Identity Optimization. 210 Supported NIC firmware versions for 10 Identity Optimization. 212 Virtual or Remote assigned Address and Persistence Policy behavior when IDRAC is set to 212 System behavior for FlexAddress and or Console mode. 213 Enabling or disabiling 10 Identity Optimization. 214 Configuring persistence policy settings. 215 Chapter 17: Managing storage devices 219 Understanding RAID concepts. 220 What is RAID. 221 Organizing data storage for availability and performance. 222 Consing RAID levels 222 Comparing RAID level performance. 229 Supported controllers. 230 Supported controllers. 230 Supported controllers. 230 Inventoring storage devices using web interface. 236 Monitoring storage devices using RACADM. 236 Monitoring storage devices using RACADM. 236 Viewing storage device topology. 236 Managing unuasigning physical disk	Serial Data Capture	
Supported cards for IO Identity Optimization. 210 Supported NIC firmware versions for IO Identity Optimization. 212 Virtual or Remote assigned Address and Persistence Policy behavior when IDRAC is set to Remote-Assigned Address mode or Console mode. 212 System behavior for FlexAddress and IO Identity. 213 Enabling or disabiling IO Identity Optimization. 214 SDB Wear Threshold. 214 Configuring persistence policy settings. 215 Chapter 17: Managing storage devices 219 Understanding RAID concepts. 220 What is RAID. 221 Organizing data storage for availability and performance. 222 Choosing RAID level performance. 228 Supported controllers. 229 Supported controllers. 230 Inventorying and monitoring storage devices. 235 Monitoring storage devices using RAD Mexices using RAD Mexices using RAD Mexices 236 Monitoring storage devices using RADA Mexices using web interface. 236 Monitoring storage devices using RADA Mexices using RADA Mexices using RADA Mexices using RADA Mexices using RADAM 236 Monitoring storage devices using RADAM 236 Monitoring storage devices using RADAM	Dynamic configuration of virtual addresses, initiator, and storage target settings	210
Supported NIC firmware versions for IO Identity Optimization. 212 Virtual or Remote assigned Address and Persistence Policy behavior when IDRAC is set to 212 System behavior for FlexAddress and IO Identity. 213 Enabling or disabling IO Identity Optimization. 214 SSD Wear Threshold. 214 Configuring persistence policy settings. 215 Chapter 17: Managing storage devices 219 Understanding RAID concepts. 220 What is RAID. 221 Organizing data storage for availability and performance. 222 Comparing RAID levels 222 Comparing RAID level performance. 228 Supported enclosures. 230 Supported controllers. 230 Supported enclosures. 230 Numery of supported features for storage devices. 235 Monitoring storage devices using RACDM. 236 Monitoring backplane using IDRAC settings utility. 236 Viewing storage device topology. 236 Managing physical disk. 237 Converting a physical disk. 238 Prasing Sto/SE device data. 239 Rebuild Physical disk. <td>Supported cards for IO Identity Optimization</td> <td>210</td>	Supported cards for IO Identity Optimization	210
Virtual or Remote assigned Address and Persistence Policy behavior when IDRAC is set to Remote-Assigned Address and IO Identity. 213 System behavior for FlexAddress and IO Identity. 213 Support of FlexAddress and IO Identity. 214 SD Wear Threshold. 214 Configuring persistence policy settings. 219 Understanding RAID concepts. 220 What is RAID. 221 Organizing data storage for availability and performance. 222 Comparing RAID levels 222 Comparing RAID level performance. 228 Supported controllers. 229 Supported controllers. 230 Supported controllers. 230 Supported controllers. 230 Supported controllers. 230 Supmetry of supported features for storage devices. 235 Monitoring storage devices using Web Interface. 236 Monitoring storage devices using RACADM. 236 Viewing storage device topology. 236 Managing physical disk. 237 Converting a physical disk. 237 Converting a physical disk. 237 Converting	Supported NIC firmware versions for IO Identity Optimization	212
System behavior for FlexAddress and IO Identity. 213 Enabling or disabling IO Identity Optimization. 214 SSD Wear Threshold. 214 Configuring persistence policy settings. 215 Chapter 17: Managing storage devices 219 Understanding RAID concepts. 220 What is RAID. 221 Organizing data storage for availability and performance. 222 Choosing RAID level performance. 222 Comparing RAID level performance. 228 Supported controllers. 229 Supported enclosures. 230 Summary of supported features for storage devices. 230 Inventorying and monitoring storage devices using RACADM. 236 Monitoring storage devices using RACADM. 236 Monitoring backplane using IDRAC settings utility. 236 Viewing storage device topology. 236 Managing physical disk to RAID or non-RAID mode. 238 Erasing physical disk to RAID or non-RAID mode. 239 Erasing physical disk to consistency. 244 Checking virtual disks. 244 Checking virtual disks. 244 Checking virtual disks.<	Virtual or Remote assigned Address and Persistence Policy behavior when iDRAC is set to Remote-Assigned Address mode or Console mode	
Enabling or disabling IO Identity Optimization. 214 SSD Wear Threshold. 214 Configuring persistence policy settings. 215 Chapter 17: Managing storage devices 219 Understanding RAID concepts. 220 What is RAID. 221 Organizing data storage for availability and performance. 222 Choosing RAID levels 222 Comparing RAID level performance. 228 Supported controllers. 229 Supported controllers. 229 Supported controllers. 229 Supported controllers. 229 Supported controllers. 230 Inventorying and monitoring storage devices 235 Monitoring storage devices using RACADM. 236 Viewing storage devices topology. 236 Managing physical disk as global hot spare. 237 Assigning or unassigning physical disk as global hot spare. 237 Assigning or unassigning physical disk as global hot spare. 238 Frasing SED/ISE device data. 239 Erasing virtual disks. 241 Creating virtual disks. 244 Checking virtual dis	System behavior for FlexAddress and IO Identity	213
SSD Wear Threshold. 214 Configuring persistence policy settings. 215 Chapter 17: Managing storage devices 219 Understanding RAID concepts. 220 What is RAID. 221 Orgenizing data storage for availability and performance. 222 Choosing RAID levels 222 Comparing RAID level performance. 228 Supported controllers. 229 Supported controllers. 230 Supmary of supported features for storage devices. 230 Inventorying and monitoring storage devices. 230 Monitoring storage devices using web interface. 235 Monitoring storage devices using KACADM. 236 Monitoring backplane using IDRAC settings utility. 236 Viewing storage device topology. 237 Assigning or unassigning physical disk as global hot spare. 237 Converting a physical disk. 239 Erasing Physical disk. 241 Creating virtual disks. 241 Changing virtual disks. 241 Changing virtual disks. 241 Changing virtual disks. 243 Deleting virtual dis	Enabling or disabling IO Identity Optimization	
Configuring persistence policy settings.215Chapter 17: Managing storage devices219Understanding RAID concepts.220What is RAID221Organizing data storage for availability and performance.222Choosing RAID levels222Comparing RAID levels222Comparing RAID levels228Supported controllers.229Supported controllers.220Supported controllers.230Summary of supported features for storage devices.230Inventorying and monitoring storage devices using RACADM.236Monitoring storage devices using RACADM.236Monitoring backplane using IDRAC settings utility.236Managing physical disks.237Assigning or unassigning physical disk as global hot spare.237Converting a physical disk to RAID or non-RAID mode.238Erasing SED/ISE device data.239Erasing SED/ISE device data.239Rebuild Physical Disk.241Checking virtual disks.241Checking virtual disks.244Checking virtual disks.244Checking virtual disks.244Checking virtual disks.245Assigning or unassigning dedicated hot spares.246Managing virtual disks.244Checking virtual disks.244Checking virtual disks.244Checking virtual disks.244Checking virtual disks.245Assigning or unassigning dedicated hot spares.245Managin	SSD Wear Threshold	
Chapter 17: Managing storage devices 219 Understanding RAID concepts. 220 What is RAID 221 Organizing data storage for availability and performance. 222 Choosing RAID levels 222 Comparing RAID levels 222 Comparing RAID level performance. 228 Supported controllers. 229 Supported controllers. 230 Summary of supported features for storage devices. 230 Inventorying and monitoring storage devices using web interface. 236 Monitoring storage devices using RACADM. 236 Monitoring storage devices using RACADM. 236 Monitoring backplane using IDRAC settings utility. 236 Viewing storage device topology. 236 Managing physical disks. 237 Assigning or unassigning physical disk as global hot spare. 237 Converting a physical disks. 239 Erasing SED/ISE device data. 239 Rebuild Physical Disk. 241 Creating virtual disks. 241 Creating virtual disks. 244 Checking virtual disks. 244 Checking virtual	Configuring persistence policy settings	215
Phapter 17: Managing storage devices 219 Understanding RAID concepts. 220 What is RAID 221 Organizing data storage for availability and performance. 222 Choosing RAID levels 222 Comparing RAID levels 222 Comparing RAID levels 223 Supported controllers. 229 Supported enclosures. 230 Summary of supported features for storage devices. 230 Inventorying and monitoring storage devices using web interface. 236 Monitoring storage devices using RACADM. 236 Monitoring backplane using iDRAC settings utility. 236 Viewing storage device topology. 236 Managing physical disk. 237 Assigning or unassigning physical disk as global hot spare. 237 Converting a physical disk. 239 Erasing SED/ISE device data. 239 Erasing physical disk. 241 Checking virtual disks. 241 Managing virtual disk. 241 Editing virtual disk. 243 Deleting virtual disk. 244 Checking virtual disk. 244 <		
Orderstanding KAID concepts. 220 What is RAID. 221 Organizing data storage for availability and performance. 222 Choosing RAID levels 222 Comparing RAID levels 222 Comparing RAID levels 222 Supported controllers. 229 Supported enclosures. 230 Summary of supported features for storage devices. 235 Monitoring storage devices using web interface. 236 Monitoring storage devices using RACADM. 236 Monitoring backplane using IDRAC settings utility. 236 Viewing storage device topology. 236 Managing physical disk. 237 Assigning or unassigning physical disk as global hot spare. 237 Converting a physical disk. 239 Erasing SED/ISE device data. 239 Rebuild Physical Disk. 241 Managing virtual disk. 241 Checking virtual disk. 241 Checking virtual disk. 241 Editing virtual disk. 244 Checking virtual disk. 245 Assigning or unassigning dedicated hot spares. 245	Chapter 17: Managing storage devices	219
What is KAD.221Organizing data storage for availability and performance.222Choosing RAID levels222Comparing RAID level performance.228Supported controllers.229Supported controllers.230Summary of supported features for storage devices.230Inventorying and monitoring storage devices.236Monitoring storage devices using web interface.236Monitoring storage devices using RACADM.236Monitoring storage devices using IDRAC settings utility.236Viewing storage device topology.236Managing physical disks.237Converting a physical disk to RAID or non-RAID mode.238Erasing SED/ISE device data.239Rebuild Physical Disk.241Managing virtual disks.241Creating virtual disks.241Creating virtual disks.241Editing virtual disks.244Checking virtual disks.245Assigning or unassigning RACADM.246Managing virtual disks using RACADM.246Managing virtual disks.247Managing virtual disks.244Checking virtual disks.245Assigning or unassigning Accated hot spares.245Managing virtual disks using RACADM.2	Understanding RAID concepts	
Organizing data storage for availability and performance. 222 Choosing RAID levels 222 Comparing RAID levels 228 Supported controllers 229 Supported enclosures 230 Summary of supported features for storage devices 230 Inventorying and monitoring storage devices 236 Monitoring storage devices using web interface 236 Monitoring storage devices using RACADM. 236 Monitoring backplane using IDRAC settings utility. 236 Managing physical disks. 237 Assigning or unassigning physical disk as global hot spare. 237 Converting a physical disks. 239 Erasing SED/ISE device data. 239 Rebuild Physical Disk. 241 Managing virtual disks. 241 Creating virtual disk. 241 Managing virtual disk. 244 Checking virtual disks. 244 Checking virtual disks. 244 Checking virtual disks. 244 Assigning or unassigning dedicated hot spares 245 Assigning or unassigning dedicated hot spares 244 Checking virtual di	What is RAID	
Choosing RAID levels222Comparing RAID level performance228Supported controllers.229Supported enclosures.230Summary of supported features for storage devices.230Inventorying and monitoring storage devices.235Monitoring storage devices using web interface.236Monitoring backplane using iDRAC settings utility.236Monitoring backplane using iDRAC settings utility.236Managing physical disks.237Assigning or unassigning physical disk as global hot spare.237Converting a physical disk to RAID or non-RAID mode.238Erasing SED/ISE device data239Rebuild Physical Disk.241Managing virtual disks.241Creating virtual disks.241Checking virtual disks.243Deleting virtual disks.244Checking virtual disks.244Checking virtual disks.244Rasigning or unassigning dedicated hot spares.243Rebuild Physical disk consistency244Checking virtual disks.244Checking virtual disks.244Checking virtual disks.245Managing virtual disks.245Managing virtual disks.246Assigning or unassigning dedicated hot spares.245Managing virtual disks using RACADM.248RAID Configuration Features.249Managing controller properties.250Configuring controller properties.250Configuring controller properties.	Organizing data storage for availability and performance	
Comparing RAID level performance228Supported controllers.229Supported enclosures.230Summary of supported features for storage devices.230Inventorying and monitoring storage devices.235Monitoring storage devices using web interface.236Monitoring storage devices using RACADM.236Monitoring backplane using iDRAC settings utility.236Viewing storage device topology.236Managing physical disks.237Converting a physical disk as global hot spare.237Converting a physical disk to RAID or non-RAID mode.238Erasing SED/ISE device data239Erasing SED/ISE device data239Rebuild Physical Disk.241Managing virtual disks.241Creating virtual disk.244Checking virtual disk.244Checking virtual disks.244Checking virtual disks.244Loncypting virtual disks.244Assigning or unassigning dedicated hot spares.245Assigning virtual disks.244Creating virtual disks.244Creating virtual disks.244Checking virtual disks.245Assigning or unassigning dedicated hot spares.245Managing virtual disks using RACADM.248RAID Configuration Features.249Managing virtual disks using RACADM.248RAID Configuration Features.240Managing virtual disks using RACADM.248RAID Configuration Features.240<	Choosing RAID levels	
Supported controllers229Supported enclosures230Summary of supported features for storage devices230Inventorying and monitoring storage devices235Monitoring storage devices using web interface236Monitoring storage devices using RACADM.236Monitoring backplane using iDRAC settings utility.236Viewing storage device topology.236Managing physical disks237Assigning or unassigning physical disk as global hot spare.238Erasing SED/ISE device data.239Erasing SED/ISE device data.239Rebuild Physical Disk241Managing virtual disks241Creating virtual disk243Deleting virtual disk244Checking virtual disks244Checking virtual disks244Checking virtual disks244Rasigning or unassigning dedicated hot spares245Managing virtual disks244Checking virtual disks244Checking virtual disks244Checking virtual disks244RAID Configuration Features245Managing virtual disks using web interface247Managing virtual disks using RACADM248RAID Configuration Features240Managing controllers250Configuring controllers250Configuring controllers250Configuring controllers250Immoduler253Managing controllers250Configuring controllers250 <td>Comparing RAID level performance</td> <td>228</td>	Comparing RAID level performance	228
Supported enclosures.230Summary of supported features for storage devices.230Inventorying and monitoring storage devices using web interface.236Monitoring storage devices using RACADM.236Monitoring backplane using iDRAC settings utility.236Viewing storage device topology.236Managing physical disks.237Assigning or unassigning physical disk as global hot spare.237Converting a physical disk to RAID or non-RAID mode.238Erasing physical disks.239Erasing physical disks.231Converting a physical Disk.241Managing virtual disks.241Creating virtual disks.241Creating virtual disks.241Creating virtual disks.241Editing virtual disk cache policies.243Deleting virtual disk cache policies.244Line virtual disk cache policies.244Checking virtual disk suing web interface.245Managing virtual disks.244Encrypting virtual disks.245Managing virtual disks using RACADM.248RAID Configuration Features.245Managing ontrollers.246Managing ontrollers.250Configuration freetor.250Importing or unastigning dedicated hot spares.250Configuration freetor.250Importing or unastigning negative configuration255	Supported controllers	
Summary of supported features for storage devices230Inventorying and monitoring storage devices235Monitoring storage devices using web interface236Monitoring storage devices using RACADM.236Monitoring backplane using iDRAC settings utility.236Viewing storage device topology.236Managing physical disks.237Assigning or unassigning physical disk as global hot spare.237Converting a physical disk to RAID or non-RAID mode.238Erasing physical disks.239Rebuild Physical Disk.241Managing virtual disks.241Creating virtual disks.241Creating virtual disks.241Checking virtual disks.241Liting virtual disks.243Deleting virtual disks.244Checking virtual disks.244Liting virtual disks.244Rebuild Physical disks.244Rebuild pirtual disks.244Checking virtual disks.244Checking virtual disks.244Checking virtual disks.245Assigning or unassigning dedicated hot spares.245Managing virtual disks using web interface.247Managing virtual disks using RACADM.248RAID Configuration Features.245Managing virtual disk using RACADM.248RAID Configuration Features.250Configuring controller properties.250Configuring controller properties.250Configuring controller properties.250<	Supported enclosures	230
Inventorying and monitoring storage devices235Monitoring storage devices using web interface.236Monitoring storage devices using RACADM.236Monitoring backplane using iDRAC settings utility.236Viewing storage device topology.236Managing physical disks.237Assigning or unassigning physical disk as global hot spare.237Converting a physical disk to RAID or non-RAID mode.238Erasing SED/ISE device data.239Rebuild Physical Disk.241Managing virtual disks.241Creating virtual disks.243Deleting virtual disks.244Checking virtual disks.244Checking virtual disks.244Encrypting virtual disks.244Rebuild Physical disks.244Rebuild pixtual disks.244Rebuild pixtual disks.244Checking virtual disks.244Checking virtual disks.244Rebuild pixtual disks.244Rebuild pixtual disks.244RASIGN or unassigning dedicated hot spares.245Assigning or unassigning dedicated hot spares.245Managing virtual disks using RACADM.248RAID Configuration Features.249Managing controllers.250Configuring controller properties.250Immorphile properties.250Immorphile properties.250Immorphile properties.250	Summary of supported features for storage devices	
Monitoring storage devices using web interface. 236 Monitoring storage devices using RACADM. 236 Monitoring backplane using iDRAC settings utility. 236 Viewing storage device topology. 236 Managing physical disks. 237 Assigning or unassigning physical disk as global hot spare. 237 Converting a physical disk to RAID or non-RAID mode. 238 Erasing physical disks. 239 Erasing SED/ISE device data. 239 Rebuild Physical Disk. 241 Managing virtual disks. 241 Creating virtual disks. 241 Checking virtual disks. 243 Deleting virtual disk consistency. 244 Checking virtual disks. 244 Checking virtual disks. 244 Checking virtual disks. 244 Checking virtual disks. 244 Managing virtual disks. 244 Managing virtual disks. 244 Checking virtual disks. 244 Checking virtual disks. 244 Managing virtual disks. 245 Assigning or unassigning dedicated hot spares. 245 <	Inventorying and monitoring storage devices	
Monitoring storage devices using RACADM.236Monitoring backplane using iDRAC settings utility.236Viewing storage device topology.236Managing physical disks.237Assigning or unassigning physical disk as global hot spare.237Converting a physical disk to RAID or non-RAID mode.238Erasing physical disks.239Erasing SED/ISE device data.239Rebuild Physical Disk.241Managing virtual disks.241Creating virtual disks.241Checking virtual disks.243Deleting virtual disks.244Checking virtual disks.244Checking virtual disks.244Lintializing virtual disks.244Assigning or unassigning dedicated hot spares.245Managing virtual disks using RACADM.248RAID Configuration Features.249Managing virtual disks using RACADM.249Managing controllers.250Configuring controller properties.250Configuring controller properties.250Importing of auto importing foreing continueration.253	Monitoring storage devices using web interface	
Monitoring backplane using iDRAC settings utility.236Viewing storage device topology.236Managing physical disks.237Assigning or unassigning physical disk as global hot spare.237Converting a physical disk to RAID or non-RAID mode.238Erasing physical disks.239Erasing SED/ISE device data.239Rebuild Physical Disk.241Managing virtual disks.241Creating virtual disks.241Creating virtual disks.243Deleting virtual disk cache policies.243Deleting virtual disk consistency.244Checking virtual disks.244Encrypting virtual disks.244Encrypting virtual disks.244Rasigning or unassigning dedicated hot spares.245Assigning or unassigning dedicated hot spares.245Managing virtual disks using web interface.247Managing virtual disks using RACADM.248RAID Configuration Features.249Managing controllers.250Contiguring controller properties.250Configuring controller properties.250Configuring controller properties.250Configuring controller properties.253	Monitoring storage devices using RACADM	
Viewing storage device topology.236Managing physical disks.237Assigning or unassigning physical disk as global hot spare.237Converting a physical disk to RAID or non-RAID mode.238Erasing physical disks.239Erasing SED/ISE device data.239Rebuild Physical Disk.241Managing virtual disks.241Creating virtual disks.241Creating virtual disks.243Deleting virtual disk cache policies.243Deleting virtual disks.244Checking virtual disks.244Locking virtual disks.244Locking virtual disks.244Locking virtual disks.244Encrypting virtual disks.244Encrypting virtual disks.245Assigning or unassigning dedicated hot spares.245Managing virtual disks using web interface.247Managing virtual disks using RACADM.248RAID Configuration Features.249Managing controllers.250Configuring controller properties.250Importing controller properties.250Importing controller properties.253	Monitoring backplane using iDRAC settings utility	
Managing physical disks237Assigning or unassigning physical disk as global hot spare237Converting a physical disk to RAID or non-RAID mode238Erasing physical disks239Erasing SED/ISE device data239Rebuild Physical Disk241Managing virtual disks241Creating virtual disks241Creating virtual disks243Deleting virtual disk243Deleting virtual disk244Checking virtual disks244Checking virtual disks244Initializing virtual disks244Encrypting virtual disks245Assigning or unassigning dedicated hot spares245Managing virtual disks using web interface247Managing virtual disks using RACADM248RAID Configuration Features249Managing controllers250Configuring controller properties250Importing or a uto importing foreign configuration253	Viewing storage device topology	
Assigning or unassigning physical disk as global hot spare.237Converting a physical disk to RAID or non-RAID mode.238Erasing physical disks.239Erasing SED/ISE device data.239Rebuild Physical Disk.241Managing virtual disks.241Creating virtual disks.241Editing virtual disks.243Deleting virtual disk cache policies.243Deleting virtual disk consistency.244Initializing virtual disks.244Encrypting virtual disks.244Assigning or unassigning dedicated hot spares.245Managing virtual disks using web interface.247Managing virtual disks using RACADM.248RAID Configuration Features.249Managing controller properties.250Configuring controller properties.250Importing or auto importing foreign configuration253	Managing physical disks	
Converting a physical disk to RAID or non-RAID mode.238Erasing physical disks.239Erasing SED/ISE device data.239Rebuild Physical Disk.241Managing virtual disks.241Creating virtual disks.241Creating virtual disks.241Editing virtual disk cache policies.243Deleting virtual disk consistency.244Checking virtual disks.244Initializing virtual disks.244Encrypting virtual disks.245Assigning or unassigning dedicated hot spares.245Managing virtual disks using web interface.247Managing virtual disks using RACADM.248RAID Configuration Features.250Configuring controller properties.250Importing or auto importing foreign configuration253	Assigning or unassigning physical disk as global hot spare	
Erasing physical disks239Erasing SED/ISE device data.239Rebuild Physical Disk241Managing virtual disks241Creating virtual disks241Creating virtual disk cache policies.243Deleting virtual disk consistency.244Checking virtual disks244Initializing virtual disks244Encrypting virtual disks245Assigning or unassigning dedicated hot spares.245Managing virtual disks using web interface.247Managing virtual disks using RACADM.248RAID Configuration Features.240Managing controllers.250Configuring controller properties.250Importing or auto importing foreign configuration253	Converting a physical disk to RAID or non-RAID mode	
Erasing SED/ISE device data.239Rebuild Physical Disk.241Managing virtual disks.241Creating virtual disks.241Editing virtual disk cache policies.243Deleting virtual disk consistency.244Checking virtual disks.244Initializing virtual disks.244Encrypting virtual disks.245Assigning or unassigning dedicated hot spares.245Managing virtual disks using RACADM.248RAID Configuration Features.249Managing controllers.250Configuring controller properties.250Importing or auto importing foreign configuration253	Erasing physical disks	239
Rebuild Physical Disk.241Managing virtual disks.241Creating virtual disks.241Editing virtual disk cache policies.243Deleting virtual disks.244Checking virtual disk consistency.244Initializing virtual disks.244Encrypting virtual disks.245Assigning or unassigning dedicated hot spares.245Managing virtual disks using web interface.247Managing virtual disks using RACADM.248RAID Configuration Features.250Configuring controllers.250Importing or auto importing foreign configuration253	Erasing SED/ISE device data	239
Managing virtual disks.241Creating virtual disks.241Editing virtual disk cache policies.243Deleting virtual disks.244Checking virtual disk consistency.244Initializing virtual disks.244Encrypting virtual disks.244Encrypting virtual disks.245Assigning or unassigning dedicated hot spares.245Managing virtual disks using web interface.247Managing virtual disks using RACADM.248RAID Configuration Features.249Managing controllers.250Configuring controller properties.250Importing or auto importing foreign configuration253	Rebuild Physical Disk	241
Creating virtual disks.241Editing virtual disk cache policies.243Deleting virtual disk consistency.244Checking virtual disk consistency.244Initializing virtual disks.244Encrypting virtual disks.245Assigning or unassigning dedicated hot spares.245Managing virtual disks using web interface.247Managing virtual disks using RACADM.248RAID Configuration Features.249Managing controllers.250Configuring controller properties.250Importing or auto importing foreign configuration.253	Managing virtual disks	
Editing virtual disk cache policies.243Deleting virtual disks.244Checking virtual disk consistency.244Initializing virtual disks.244Encrypting virtual disks.245Assigning or unassigning dedicated hot spares.245Managing virtual disks using web interface.247Managing virtual disks using RACADM.248RAID Configuration Features.249Managing controllers.250Configuring controller properties.250Importing or auto importing foreign configuration253	Creating virtual disks	241
Deleting virtual disks244Checking virtual disk consistency244Initializing virtual disks244Encrypting virtual disks245Assigning or unassigning dedicated hot spares245Managing virtual disks using web interface247Managing virtual disks using RACADM248RAID Configuration Features249Managing controllers250Configuring controller properties250Importing or auto importing foreign configuration253	Editing virtual disk cache policies	243
Checking virtual disk consistency.244Initializing virtual disks.244Encrypting virtual disks.245Assigning or unassigning dedicated hot spares.245Managing virtual disks using web interface.247Managing virtual disks using RACADM.248RAID Configuration Features.249Managing controllers.250Configuring controller properties.250Importing or auto importing foreign configuration253	Deleting virtual disks	244
Initializing virtual disks244Encrypting virtual disks245Assigning or unassigning dedicated hot spares245Managing virtual disks using web interface247Managing virtual disks using RACADM248RAID Configuration Features249Managing controllers250Configuring controller properties250Importing or auto importing foreign configuration253	Checking virtual disk consistency	
Encrypting virtual disks	Initializing virtual disks	244
Assigning or unassigning dedicated hot spares	Encrypting virtual disks	
Managing virtual disks using web interface. 247 Managing virtual disks using RACADM. 248 RAID Configuration Features. 249 Managing controllers. 250 Configuring controller properties. 250 Importing or auto importing foreign configuration 253	Assigning or unassigning dedicated hot spares	
Managing virtual disks using RACADM	Managing virtual disks using web interface	
RAID Configuration Features	Managing virtual disks using RACADM	
Managing controllers	RAID Configuration Features	
Configuring controller properties	Managing controllers	250
Importing or auto importing foreign configuration	Configuring controller properties	250
	Importing or auto importing foreign configuration	255

OTOCO

10 Contents

	In In
Clearing foreign configuration	
Resetting controller configuration	
Switching the controller mode	255
12 Gbps SAS HBA adapter operations	257
Monitoring predictive failure analysis on drives	257
Controller operations in non-RAID mode or HBA mode	258
Running RAID configuration jobs on multiple storage controllers	
Manage Preserved cache	258
Managing PCIe SSDs	259
Inventorying and monitoring PCIe SSDs	259
Preparing to remove PCIe SSD	260
Erasing PCIe SSD device data	
Managing enclosures or backplanes	
Configuring backplane mode	
Viewing universal slots	
Setting SGPIO mode	
Set Enclosure Asset Tag	
Set Enclosure Asset Name	266
Choosing operation mode to apply settings	
Choosing operation mode using web interface	
Choosing operation mode using RACADM	
Viewing and applying pending operations	
Viewing, applying, or deleting pending operations using web interface	267
Viewing and applying pending operations using RACADM	268
Storage devices — apply operation scenarios	268
Blinking or unblinking component LEDs	
Blinking or unblinking component LEDs using web interface	
Blinking or unblinking component LEDs using RACADM	
Warm reboot	270
Chapter 18: BIOS Settings	271
BIOS Live Scanning	272
BIOS Recovery and Hardware Root of Trust (RoT)	273
Chapter 19: Configuring and using virtual console	274
Supported screen resolutions and refresh rates	
Configuring virtual console	
Configuring virtual console using web interface	276
Configuring virtual console using RACADM	276
Previewing virtual console	
Launching virtual console	276
Launching virtual console using web interface	277
Launching virtual console using a URL	277
Disabling warning messages while launching virtual console or virtual media using Java or Activ	veX
plug-in	
Using virtual console viewer	278
eHTML5 based virtual console	278
HTML5 based virtual console	281
Synchronizing mouse pointers	
Passing all keystrokes through virtual console for Java or ActiveX plug-in	

FIS. <u>397</u> Mov. <u>32</u>



Chapter 20: Using iDRAC Service Module	
Installing iDRAC Service Module	
Installing iDRAC Service Module from iDRAC Express and Basic	
Installing iDRAC Service Module from iDRAC Enterprise	
Supported operating systems for iDRAC Service Module	
iDRAC Service Module monitoring features	
Using iDRAC Service Module from iDRAC web interface	294
Using iDRAC Service Module from RACADM	
Chapter 21: Using USB port for server management	
Accessing iDRAC interface over direct USB connection	
Configuring iDRAC using server configuration profile on USB device	
Configuring USB management port settings	
Importing Server Configuration Profile from USB device	
Chapter 22: Using Quick Sync 2	300
Configuring iDRAC Quick Sync 2	
Configuring iDRAC Quick Sync 2 settings using web interface	
Configuring iDRAC Quick Sync 2 settings using RACADM	
Configuring iDRAC Quick Sync 2 settings using iDRAC settings utility	
Using mobile device to view iDRAC information	
Chapter 23: Managing virtual media	302
Supported drives and devices	
Configuring virtual media	
Configuring virtual media using iDRAC web interface	
Configuring virtual media using RACADM	
Configuring virtual media using iDRAC settings utility	
Attached media state and system response	
Accessing virtual media	
Launching virtual media using virtual console	
Adding virtual media without using virtual console	
Adding virtual media images	
Mapping virtual drive	
Inmapping virtual drive	307
Setting boot order through BIOS	
Enabling boot once for virtual media	
Chapter 24: Managing vElash SD card	309
Configuring vElash SD card	
Viewing vFlash SD card properties	
Enabling or disabling vFlash functionality	
Initializing vFlash SD card	
Getting the last status using RACADM	
Managing vFlash partitions	



Creating a partition using an image file	
Formatting a partition	
Viewing available partitions	
Modifying a partition	
Attaching or detaching partitions	
Deleting existing partitions	
Downloading partition contents	
Booting to a partition	
Chapter 25: Using SMCLP	
System management capabilities using SMCLP	
Running SMCLP commands	
iDRAC SMCLP syntax	
Navigating the map address space	
Using show verb	
Using the -display option	
Using the -level option	
Using the -output option	
Usage examples	
Server power management	
SEL management	
Map target navigation	
Chapter 26: Deploying operating systems	
Deploying operating system using remote file share	
Managing remote file shares	
Configuring remote file share using web interface	
Configuring remote file share using RACADM	
Deploving operating system using virtual media	
bopioying operating system abing virtual modularianianianianianianianianiani	
Installing operating system from multiple disks	
Installing operating system from multiple disks Deploying embedded operating system on SD card	
Installing operating system from multiple disks Deploying embedded operating system on SD card Enabling SD module and redundancy in BIOS	
Installing operating system from multiple disks Deploying embedded operating system on SD card Enabling SD module and redundancy in BIOS	
Installing operating system from multiple disks Deploying embedded operating system on SD card Enabling SD module and redundancy in BIOS Chapter 27: Troubleshooting managed system using iDRAC	
Installing operating system from multiple disks Deploying embedded operating system on SD card Enabling SD module and redundancy in BIOS Chapter 27: Troubleshooting managed system using iDRAC Using diagnostic console	
Installing operating system from multiple disks Deploying embedded operating system on SD card Enabling SD module and redundancy in BIOS Chapter 27: Troubleshooting managed system using iDRAC Using diagnostic console Reset iDRAC and Reset iDRAC to default	
Installing operating system from multiple disks Deploying embedded operating system on SD card Enabling SD module and redundancy in BIOS Chapter 27: Troubleshooting managed system using iDRAC Using diagnostic console Reset iDRAC and Reset iDRAC to default	
Installing operating system from multiple disks Deploying embedded operating system on SD card Enabling SD module and redundancy in BIOS Chapter 27: Troubleshooting managed system using iDRAC Using diagnostic console Reset iDRAC and Reset iDRAC to default Scheduling remote automated diagnostics Scheduling remote automated diagnostics using RACADM	
Installing operating system from multiple disks Deploying embedded operating system on SD card Enabling SD module and redundancy in BIOS Chapter 27: Troubleshooting managed system using iDRAC Using diagnostic console Reset iDRAC and Reset iDRAC to default Scheduling remote automated diagnostics Scheduling remote automated diagnostics using RACADM Viewing post codes	
Installing operating system from multiple disks Deploying embedded operating system on SD card Enabling SD module and redundancy in BIOS Chapter 27: Troubleshooting managed system using iDRAC Using diagnostic console Reset iDRAC and Reset iDRAC to default Scheduling remote automated diagnostics Scheduling remote automated diagnostics using RACADM Viewing post codes Viewing boot and crash capture videos	
Installing operating system from multiple disks Deploying embedded operating system on SD card Enabling SD module and redundancy in BIOS Chapter 27: Troubleshooting managed system using iDRAC Using diagnostic console Reset iDRAC and Reset iDRAC to default Scheduling remote automated diagnostics Scheduling remote automated diagnostics using RACADM Viewing post codes Viewing boot and crash capture videos Configuring video capture settings	
Installing operating system from multiple disks Deploying embedded operating system on SD card Enabling SD module and redundancy in BIOS Chapter 27: Troubleshooting managed system using iDRAC Using diagnostic console Reset iDRAC and Reset iDRAC to default Scheduling remote automated diagnostics Scheduling remote automated diagnostics using RACADM Viewing post codes Viewing boot and crash capture videos Configuring video capture settings Viewing logs	
Installing operating system from multiple disks Deploying embedded operating system on SD card Enabling SD module and redundancy in BIOS Chapter 27: Troubleshooting managed system using iDRAC Using diagnostic console Reset iDRAC and Reset iDRAC to default Scheduling remote automated diagnostics Scheduling remote automated diagnostics using RACADM Viewing post codes Viewing boot and crash capture videos Configuring video capture settings Viewing logs Viewing logs	
Installing operating system from multiple disks Deploying embedded operating system on SD card Enabling SD module and redundancy in BIOS Chapter 27: Troubleshooting managed system using iDRAC Using diagnostic console Reset iDRAC and Reset iDRAC to default Scheduling remote automated diagnostics Scheduling remote automated diagnostics using RACADM Viewing post codes Viewing boot and crash capture videos Configuring video capture settings Viewing logs Viewing logs Viewing last system crash screen Viewing System status	
Installing operating system from multiple disks Deploying embedded operating system on SD card Enabling SD module and redundancy in BIOS Chapter 27: Troubleshooting managed system using iDRAC Using diagnostic console Reset iDRAC and Reset iDRAC to default Scheduling remote automated diagnostics Scheduling remote automated diagnostics using RACADM Viewing post codes Viewing boot and crash capture videos Configuring video capture settings Viewing logs Viewing last system crash screen Viewing System status Viewing system front panel LCD status	
Installing operating system from multiple disks Deploying embedded operating system on SD card Enabling SD module and redundancy in BIOS Chapter 27: Troubleshooting managed system using iDRAC Using diagnostic console Reset iDRAC and Reset iDRAC to default Scheduling remote automated diagnostics Scheduling remote automated diagnostics using RACADM Viewing post codes Viewing boot and crash capture videos Configuring video capture settings Viewing logs Viewing last system crash screen Viewing System status Viewing system front panel LCD status Viewing system front panel LED status	
Installing operating system from multiple disks Deploying embedded operating system on SD card Enabling SD module and redundancy in BIOS Chapter 27: Troubleshooting managed system using iDRAC Using diagnostic console Reset iDRAC and Reset iDRAC to default Scheduling remote automated diagnostics Scheduling remote automated diagnostics using RACADM Viewing post codes Viewing boot and crash capture videos Configuring video capture settings Viewing logs Viewing last system crash screen Viewing System status Viewing system front panel LCD status Viewing system front panel LED status Hardware trouble indicators	
Installing operating system from multiple disks Deploying embedded operating system on SD card Enabling SD module and redundancy in BIOS Chapter 27: Troubleshooting managed system using iDRAC Using diagnostic console Reset iDRAC and Reset iDRAC to default Scheduling remote automated diagnostics Scheduling remote automated diagnostics using RACADM Viewing post codes Viewing boot and crash capture videos Configuring video capture settings Viewing logs Viewing logs Viewing logs Viewing system crash screen Viewing system front panel LCD status Viewing system front panel LED status Viewing system health	



	\ \
Checking server status screen for error messages	
Restarting iDRAC	
Reset to Custom Defaults (RTD)	
Resetting iDRAC using iDRAC web interface	
Resetting iDRAC using RACADM	
Erasing system and user data	
Resetting iDRAC to factory default settings	
Resetting iDRAC to factory default settings using iDRAC web interface	
Resetting iDRAC to factory default settings using iDRAC settings utility	
Chapter 28: SupportAssist Integration in iDRAC	
SupportAssist Registration	
Installing Service Module	
Server OS Proxy Information	
SupportAssist	
Service Request Portal	
Collection Log	
Generating SupportAssist Collection	
Generating SupportAssist Collection manually using iDRAC web interface	
Settings	
Collection Settings	
Contact Information	
Chapter 29: Frequently asked questions	
System Event Log	
Custom sender email configuration for iDRAC alerts	

Custom sender email configuration for iDRAC alerts	
Network security	
Telemetry streaming	
Active Directory	
Single Sign-On	
Smart card login	
Virtual console	
Virtual media	
vFlash SD card	
SNMP authentication	
Storage devices	
GPU (Accelerators)	
iDRAC Service Module	
RACADM	
Permanently setting the default password to calvin	
Miscellaneous	

Chapter 30: Use case scenarios	
· Troubleshooting an inaccessible managed system	
Obtaining system information and assess system health	
Setting up alerts and configuring email alerts	
Viewing and exporting System Event Log and Lifecycle Log	
Interfaces to update iDRAC firmware	
Performing graceful shutdown	



Creating new administrator user account	.362
Launching servers remote console and mounting a USB drive	363
Installing bare metal OS using attached virtual media and remote file share	363
Managing rack density	.363
Installing new electronic license	.363
Applying IO Identity configuration settings for multiple network cards in single host system reboot	363



Overview of iDRAC

The Integrated Dell Remote Access Controller (iDRAC) is designed to make you more productive as a system administrator and improve the overall availability of Dell EMC servers. iDRAC alerts you to system issues, helps you to perform remote management, and reduces the need for physical access to the system.

iDRAC technology is part of a larger data center solution that increases availability of business critical applications and workloads. The technology allows you to deploy, monitor, manage, configure, update, and troubleshoot Dell EMC systems from any location without using any agents or an operating system.

Several products work with the iDRAC to simplify and streamline IT operations. Following are some of the tools:

- OpenManage Enterprise
- OpenManage Power Center Plug in
- OpenManage Integration for VMware vCenter
- Dell Repository Manager

iDRAC is available in the following variants:

- iDRAC Basic Available by default for 100-500 series servers
- iDRAC Express Available by default on all 600 and higher series of rack or tower servers, and all blade servers
- iDRAC Enterprise Available on all server models
- iDRAC Datacenter Available on all server models

Topics:

- Benefits of using iDRAC
- Key features
- New features added
- How to use this guide
- Supported web browsers
- iDRAC licenses
- Licensed features in iDRAC9
- Interfaces and protocols to access iDRAC
- iDRAC port information
- Other documents you may need
- Contacting Dell
- Accessing documents from Dell support site
- Accessing Redfish API Guide

Benefits of using iDRAC

The benefits include:

- Increased Availability Early notification of potential or actual failures that help prevent a server failure or reduce recovery time after failure.
- Improved Productivity and Lower Total Cost of Ownership (TCO) Extending the reach of administrators to larger numbers of distant servers can make IT staff more productive while driving down operational costs such as travel.
- Secure Environment By providing secure access to remote servers, administrators can perform critical management functions while maintaining server and network security.
- Enhanced Embedded Management through Lifecycle Controller Lifecycle Controller provides deployment and simplified serviceability through Lifecycle Controller GUI for local deployment and Remote Services (WSMan) interfaces for remote deployment integrated with Dell OpenManage Enterprise and partner consoles.

For more information about Lifecycle Controller GUI, see *Lifecycle Controller User's Guide* and for remote services, see *Lifecycle Controller Remote Services Quick Start Guide* available at https://www.dell.com/idracmanuals.



Key features

The key features of iDRAC include:

NOTE: Some features are available only with iDRAC Enterprise or Datacenter license. For information on the features available for a license, see iDRAC licenses .

Inventory and Monitoring

- Telemetry data streaming.
- View managed server health.
- Inventory and monitor network adapters and storage subsystem (PERC and direct attached storage) without any operating system agents.
- View and export system inventory.
- View sensor information such as temperature, voltage, and intrusion.
- Monitor CPU state, processor automatic throttling, and predictive failure.
- View memory information.
- Monitor and control power usage.
- Support for SNMPv3 gets and alerts.
- For blade servers: launch Management Module web interface, view OpenManage Enterprise (OME) Modular information, and WWN/MAC addresses.

NOTE: CMC provides access to iDRAC through the M1000E Chassis LCD panel and local console connections. For more information, see *Chassis Management Controller User's Guide* available at https://www.dell.com/cmcmanuals.

- View network interfaces available on host operating systems.
- iDRAC9 provides improved monitoring and management functionality with Quick Sync 2. You need OpenManage Mobile app configured in your Android or iOS mobile device.

Deployment

- Manage vFlash SD card partitions.
- Configure front panel display settings.
- Manage iDRAC network settings.
- Configure and use virtual console and virtual media.
- Deploy operating systems using remote file share, and virtual media.
- Enable auto-discovery.
- Perform server configuration using the export or import XML or JSON profile feature through RACADM, WSMan
- and Redfish. For more information, see *Lifecycle Controller Remote Services Quick Start Guide* available at https://www.dell.com/idracmanuals.
- Configure persistence policy for virtual addresses, initiator, and storage targets.
- Remotely configure storage devices attached to the system at run-time.
- Perform the following operations for storage devices:
 - Physical disks: Assign or unassign physical disk as a global hot spare.
 - Virtual disks:
 - Create virtual disks.
 - Edit virtual disks cache policies.
 - Check virtual disk consistency.
 - Initialize virtual disks.
 - Encrypt virtual disks.
 - Assign or unassign dedicated hot spare.
 - Delete virtual disks.
 - Controllers:
 - Configure controller properties.
 - Import or auto-import foreign configuration.
 - Clear foreign configuration.
 - Reset controller configuration.
 - Create or change security keys.
 - PCIe SSD devices:
 - Inventory and remotely monitor the health of PCIe SSD devices in the server.
 - Prepare the PCIe SSD to be removed.
 - Securely erase the data.



- Set the backplane mode (unified or split mode).
- Blink or unblink component LEDs.
- Apply the device settings immediately, at next system reboot, at a scheduled time, or as a pending operation to be applied as a batch as part of the single job.

Update

- Manage iDRAC licenses.
- Update BIOS and device firmware for devices supported by Lifecycle Controller.
- Update or rollback iDRAC firmware and Lifecycle Controller firmware using a single firmware image.
- Manage staged updates.
- Access iDRAC interface over direct USB connection.
- Configure iDRAC using Server Configuration Profiles on USB device.

Maintenance and Troubleshooting

- Perform power-related operations and monitor power consumption.
- Optimize system performance and power consumption by modifying the thermal settings.
- No dependency on OpenManage Server Administrator for generation of alerts.
- Log event data: Lifecycle and RAC logs.
- Set email alerts, IPMI alerts, remote system logs, WS Eventing logs, Redfish event, and SNMP traps (v1, v2c, and v3) for events and improved email alert notification.
- Capture last system crash image.
- View boot and crash capture videos.
- Out-of-band monitor and alert the performance index of CPU, memory, and I/O modules.
- Configure warning threshold for inlet temperature and power consumption.
- Use iDRAC Service Module to:
 - View operating system information.
 - Replicate Lifecycle Controller logs to operating system logs.
 - Automate system recovery options.
 - Enable or disable status of Full Power Cycle for all System components except the PSU.
 - Remotely hard-reset iDRAC
 - Enable in-band iDRAC SNMP alerts
 - Access iDRAC using host OS (experimental feature)
 - Populate Windows Management Instrumentation (WMI) information.
 - Integrate with SupportAssist collection. This is applicable only if iDRAC Service Module Version 2.0 or later is installed.
 - Generate SupportAssist collection in the following ways:
 - Automatic Using iDRAC Service Module that automatically invokes the OS Collector tool.

Dell Best Practices regarding iDRAC

- Dell iDRAC's are intended to be on a separate management network; they are not designed nor intended to be placed on or connected directly to the Internet. Doing so could expose the connected system to security and other risks for which Dell is not responsible.
- Dell EMC recommends using the Dedicated Gigabit Ethernet port available on rack and tower servers. This interface is not shared with the host operating system and routes the management traffic to a separate physical network, enabling it to be separated from the application traffic. This option implies that iDRAC's dedicated network port routes its traffic separately from the server's LOM or NIC ports. The Dedicated option allows iDRAC to be assigned an IP address from the same subnet or different subnet in comparison to the IP addresses assigned to the Host LOM or NICs.
- Along with locating iDRACs on a separate management subnet, users should isolate the management subnet/vLAN with technologies such as firewalls, and limit access to the subnet/vLAN to authorized server administrators.

Secure Connectivity

Securing access to critical network resources is a priority. iDRAC implements a range of security features that includes:

- Custom signing certificate for Secure Socket Layer (SSL) certificate.
- Signed firmware updates.
- User authentication through Microsoft Active Directory, generic Lightweight Directory Access Protocol (LDAP) Directory Service, or locally administered user IDs and passwords.
- Two-factor authentication using the Smart–Card logon feature. The two-factor authentication is based on the physical smart card and the smart card PIN.
- Single Sign-On and Public Key Authentication.
- Role-based authorization, to configure specific privileges for each user.



- SNMPv3 authentication for user accounts stored locally in the iDRAC. It is recommended to use this, but it is disabled by default.
- User ID and password configuration.
- Default login password modification.
- Set user passwords and BIOS passwords using one-way hash format for improved security.
- FIPS 140-2 Level 1 capability.
- Session time-out configuration (in seconds).
- Configurable IP ports (for HTTP, HTTPS, SSH, Virtual Console, and Virtual Media).
- Secure Shell (SSH) that uses an encrypted transport layer for higher security.
- Login failure limits per IP address, with login blocking from that IP address when the limit is exceeded.
- Limited IP address range for clients connecting to iDRAC.
- Dedicated Gigabit Ethernet adapter available on rack and tower servers (additional hardware may be required).

New features added

This section provides the list of new features added in the following releases:

Firmware version 4.40.00.00

This release includes all the features from the previous releases. Following are the new features that are added in this release:

NOTE: For information about supported systems, refer to the respective version of Release Notes available at https://www.dell.com/support/article/sln308699.

- Added support for enhanced HTML5 (eHTML5) virtual KVM feature in virtual console
- Added support for eHTML5 virtual media
- Enhancement in Storage GUI page
- Added support for direct updates SEP backplane
- Added support for new update for PSU update
- Added support for uploading custom defaults and reset iDRAC to default settings using custom defaults
- Enhanced system lockdown mode support for supported devices

Following are the list of other features added in this release:

- Automation
- Support for Redfish Updates
- Monitoring/Alerting/Troubleshooting
 - FPGA Monitoring
 - SMART data logs enhancements including historical recording
 - Discrete voltage sensor reporting
 - Report actual start and completion info for job queue entries which require a server reboot to apply (Example: BIOS update).
 - Providing CPU serial numbers in SupportAssist Collection
 - **Telemetry** (requires iDRAC Datacenter license)
 - Multi-client support
 - Granular metric report options
 - Provision to POST a new custom MRD (Metric Report Definition) using any of the available 193 Metric Definitions and set desired Report Interval (referred as Recurrence Interval in MRD)
 - A single MRD can have a maximum of 68 Metric Definitions (Metric IDs)
 - Provision to create up to 24 new custom MRDs which in turn will have 24 new Metric Reports. An iDRAC can support a maximum of 48 Metric Reports (24 Pre-canned and 24 Custom
- Security
 - Automatic Certificate Enrollment Enhancements (requires iDRAC Datacenter License)
 - Integrate RSA SecurID Client into iDRAC for 2FA (requires iDRAC Datacenter License)
 - Compliance with STIG requirement "network device must authenticate NTP"
 - Removal of Telnet and TLS 1.0 from web server

• Platform feature support

BOSS 1.5 updates



Infiniband support

In 4.40.00.00 release, following features are added in Storage page on iDRAC GUI:

- From the Dashboard, you can see suggested actions to solve any health alters.
- The Storage page has been modified to included tabs for storage monitoring information, a Storage Hardware and Software Inventory, a list of Pending and Current storage jobs, and SEKM.
 - From the Storage Inventory, users can find all storage related hardware and software.
 - \circ The Pending and Current Jobs tab allows users to queue and monitor jobs from a centralized location.
 - You can also configure SEKM via the Storage page.
- When monitoring storage devices, you can customize the columns that are displayed for each device table. Column customization will be saved and persist between user sessions.
- New basic and advanced filters provided on each device page allow you to easily and efficiently customize the list of objects displayed.
- The Storage Configuration wizard has two options to create a Virtual disk Basic and Advanced.
 - In the basic Virtual Disk wizard, you can quickly create a VD from a list of available RAID configurations. iDRAC will automatically set the default values of the VD to streamline the process.
 - For the Advanced Virtual Disk wizard, you can select all the details of the VD. You can create a new volume for the VD or select and existing volume.
- Each devices pages has new global actions that allow you to show related devices or preform group operations.
 - For example, you can choose physical disks and perform group operation such as Blink, Unblink, and Create Virtual Disk.
 - Also, you can view the Physical disk inventory and create a Virtual Disk by choosing the drives without having to navigate away from the screen.
- Instead of the numerical value, the size of the physical disk is shown as a data visualization with values on the scale.
- This gives you an idea of used and available space on the drive.
- You can filter disks based on the various physical disk properties.
 - The filtering properties are displayed so that the user knows what filtering is currently being applied.

Firmware version 4.30.30.30

This release includes all the features from the previous releases. Following are the new features that are added in this release:

() NOTE: For information about supported systems, refer to the respective version of Release Notes available at https://www.dell.com/support/article/sln308699.

- Added support for PERC 11 for AMD systems
- Added support for NVMe drives behind PERC 11
- Added support for HBA11 for AMD systems
- Added support for CUPS for AMD systems
- Added support for Boot Optimized storage solution 1.5 (BOSS1.5/BOSS-S2)
- Added support for BOSS 1.5 secure firmware update
- Added support for new Matrox video driver
- Added support for NVMe Opal SEDs
- Added support for HW chain of Trust Secure Boot
- Added InifiniBand Adaptor support for Mellanox CX6
- Added support for 24x NVMe backplane for PowerEdge C6525
- Added support for new Matrox video driver
- Added support for Starlord (ConnectX-6 Dx 100GbE) to iDRAC
- Added FQDD related changes for BOSS-S2/PERC 11/HBA 11
- Added support for storage devices (like M.2 and U.2 but not limited) without a backplane
- Added support for Secure Enterprise Key Management (SEKM) for NVMe drives
- Enhanced iDRAC memory from 512MB to 1024MB
- RESTUI changes for failed Email Sending when authentication is disabled



Firmware version 4.20.20.20

Following features were added in this release:

Power Supply Unit (PSU)

- Support for 1100W ~48W DC PSU.
- Removed 4S PSU restriction.

NICs

- Support (4x 10/25 SFP28) OCP 3.0 Dell part # JTK7F Broadcom.
- Support (4x10/25) MX Mezz, Dell part # DCWFP Broadcom and MX 25G Quad port on MX platform.
- Support for adding Broadcom 10GbE NIC card support to R340.

Accelerators and CPU's

- Support for 2 new GPU cards to the Precision 7920 Rack (Navi10DT/W5700, Navi14DT/W5500).
- Support for Nvidia V100S for PowerEdge.
- Support for new Intel Processors: 6250 and 6256.

NVMe

• Support for Samsung PM 1735 and PM 1733 NVMe PCIe storage.

Automation/Scripting/Telemetry

- Support for Redfish 2018R3, 2019R1 & 2019R2 Features.
- Support of CLI method to retrieve POST codes.
- Support for Report Interval limit on Telemetry CUPS to be increased from 1 minute to 1 hour in Power Manager Plug-in.
- Support for Telemetry (Metric Report Enable/ Disable).
- Support for enhanced user logging with SSH.
- Support for adding tier specification flag to PCI Add IPMI command.

Miscellaneous

- Support for Thermal Sensor Board Cable Detection when C6420 chassis with one or more sleds is powered up.
- Support for displaying Slot number in the SLED GUI for 6420.
- Support for always arm AEP and BPS memory for an ADR flow on AC loss or Global Reset.
- Support for 10x2.5" BP/chassis part number change.
- Support for enabling "Unsupported Config" to the SEL log.

Firmware version 4.10.10.10

Following features are added in this release:

Features supported with Default license

• BIOS recovery and Root of Trust (RoT)



Features supported with Enterprise license

• Secure Enterprise Key Management (SEKM) — Added support for Vormetric Data Security Manager.

Features supported with Datacenter license

• BIOS live scanning — Only for AMD systems.

Firmware version 4.00.00.00

This release includes all the features from the previous releases. Following are the new features that are added in this release:

NOTE: For information about supported systems, refer to the respective version of Release Notes available at https://www.dell.com/support/article/sln308699.

Features supported with Datacenter license

- Telemetry streaming metric reports streamed to an analytics tool
- GPU inventory and monitoring
- Thermal Manage Advanced Power and cooling features
- Auto Certificate Enrollment and renewal for SSL certificates
- Virtual Clipboard Support Cut and Paste of text strings into the remote virtual console desktop
- SFP Transceiver Input/Output monitoring
- SMART logs Storage drives
- System Serial Data Buffer Capture
- Idle Server detection

Features supported with Enterprise or Datacenter license

- Multi Factor Authentication via email
- Agent Free Crash Video Capture (Windows only)
- Connection View for LLDP transmit
- System Lockdown mode new icon in header available from any page
- Group Manager 250 node support
- Enhanced support for Secure Enterprise Key Management (SEKM)

Features supported with default license (iDRAC Basic or iDRAC Express)

GUI Enhancement

- \circ $\;$ Task Summary section in the dashboard $\;$
- Search box in the header
- SupportAssist Collection Viewer displays the output in iDRAC GUI
- API, CLI, and SCP
 - Operating system deployment by Server Configuration Profile (SCP)
 - Enable and disable boot order cotrol to SCP and RACADM
 - New schemas to Redfish APIs
 - Option to change boot source state in SCP
 - Automation for Command/attribute auto completion in RACADM

Alerts and Monitoring

- Custom Sender Email Address for email alerts in SMTP configuration
- Cloud based email server in SMTP
- o SMARTlogs in SupportAssist log collection for hard drives and PCle SSD devices
- Include Part Number of failed component in alert messages
- Security
- 22 Overview of iDRAC



- Multiple IP filtering ranges using RACADM commands only
- iDRAC user password maximum length extended to 40 characters
- SSH Public Keys through SCP
- Customizable Security banner to SSH login
- Force Change Password (FCP) for login

• Storage and Storage Controllers

• Enable PERC to switch to SEKM encryption mode

How to use this guide

The contents of this user's guide enable you to perform various tasks using:

- iDRAC web interface Only the task-related information is provided here. For information about the fields and options, see the iDRAC Online Help that you can access from the web interface.
- RACADM The RACADM command or the object that you must use is provided here. For more information, see the *iDRAC* RACADM CLI Guide available at https://www.dell.com/idracmanuals.
- iDRAC Settings Utility Only the task-related information is provided here. For information about the fields and options, see the *iDRAC Settings Utility Online Help* that you can access when you click **Help** in the iDRAC Settings GUI (press <F2> during boot, and then click **iDRAC Settings** on the **System Setup Main Menu** page).
- Redfish Only the task-related information is provided here. For information about the fields and options, see the *iDRAC Redfish API Guide* available at www.api-marketplace.com.

Supported web browsers

iDRAC is supported on the following browsers:

- Internet Explorer/Edge
- Mozilla Firefox
- Google Chrome
- Safari

For the list of supported versions, see the *iDRAC Release Notes* available at https://www.dell.com/idracmanuals.

Supported OS and Hypervisors

iDRAC is supported on the following OS, Hypervisors:

- Microsoft Windows Server and Windows PE
- VMware ESXI
- RedHat Enterprise Linux
- SuSe Linux Enterprise Server

(i) NOTE: For the list of supported versions, see the *iDRAC Release Notes* available at https://www.dell.com/idracmanuals.

iDRAC licenses

iDRAC features are available based on the type of the license. Depending on the system model, iDRAC Basic or iDRAC Express license is installed by default. iDRAC Enterprise license, iDRAC Datacenter license, and iDRAC Secure Enterprise Key Manager (SEKM) license are available as an upgrade and can be purchased anytime. Only licensed features are available in the interfaces that enable you to configure or use iDRAC. For more information, see Licensed features in iDRAC9.

Types of licenses

iDRAC Basic or iDRAC Express are the standard licenses available by default on your system. iDRAC Enterprise and Datacenter licenses includes all the licensed features and can be purchased at any time. The types of upsell offered are:

• 30-day evaluation—Evaluation licenses are duration-based and the timer runs when power is applied to the system. This license cannot be extended.



- Perpetual—The license is bound to the Service Tag and is permanent.
- Following table lists the default license available on the following systems:

iDRAC Basic License	iDRAC Express License	iDRAC Enterprise License	iDRAC Datacenter License
PowerEdge Rack/Tower severs series 100-500	 PowerEdge C41XX PowerEdge FC6XX PowerEdge R6XX PowerEdge R64XX PowerEdge R74XX PowerEdge R74XX4 PowerEdge R74XX PowerEdge R8XX PowerEdge R9XX PowerEdge R9XX PowerEdge R9XX PowerEdge T6XX Dell Precision Rack R7920 	All platforms, with upgrade option	All platforms, with upgrade option

Table 1. Default License

iDRAC Express License	iDRAC Enterprise License	iDRAC Datacenter License
PowerEdge C41XX	All platforms, with upgrade option	All platforms, with upgrade option
PowerEdge FC6XX		
PowerEdge R6XX		
PowerEdge R64XX		
PowerEdge R7XX		
PowerEdge R74XXd		
PowerEdge R74XX		
PowerEdge R8XX		
PowerEdge R9XX		
PowerEdge R9XX		
PowerEdge T6XX		
Dell Precision Rack R7920		

NOTE: The default license available with PowerEdge C64XX systems is BMC. The BMC license was custom made for C64XX systems.

(i) NOTE: Express for Blades license is the default license for PowerEdge M6XX and MXXXX systems.

Methods for acquiring licenses

Use any of the following methods to acquire the licenses:

• Dell Digital Locker — Dell Digital Locker allows you to view and manage your products, software, and licensing information in one location. A link to the Dell Digital Locker is available in DRAC web interface- go to **Configuration** > **Licenses**.

(i) NOTE: To know more about Dell Digital Locker, refer to FAQ on the website.

- Email License is attached to an email that is sent after requesting it from the technical support center.
- Point-of-sale License is acquired while placing the order for a system.

(i) NOTE: To manage licenses or purchase new licenses, go to the Dell Digital Locker.

Acquiring license key from Dell Digital Locker

To obtain the license key from your account, you must first register your product using the registration code that is sent in the order confirmation email. This code must be entered in the **Product Registration** tab after logging into Dell Digital Locker.

From the left pane, click the **Products** or **Order History** tab to view the list of your products. Subscription-based products are listed under **Billing accounts** tab.

To download the license key from your Dell Digital Locker account:

- 1. Sign in to your Dell Digital Locker account.
- 2. From the left pane, click Products.
- 3. Click the product that you want to view.
- 4. Click the product name.
- 5. On the Product management page, click Get Key.
- 6. Follow the instructions on the screen to obtain the license key.

NOTE: If you do not have a Dell Digital Locker account, create an account using the email address provided during your purchase.

NOTE: To generate multiple license keys for new purchases, follow the instructions under Tools > License Activation >
 Unactivated licenses

License operations

Before you perform the license management tasks, ensure that you acquire the licenses. For more information, see the Methods for acquiring licenses.

(i) NOTE: If you have purchased a system with all the licenses pre-installed, then license management is not required.

You can perform the following licensing operations using iDRAC, RACADM, WSMan, Redfish and Lifecycle Controller-Remote Services for one-to-one license management, and Dell License Manager for one-to-many license management:

- View View the current license information.
- Import After acquiring the license, store the license in a local storage and import it into iDRAC using one of the supported interfaces. The license is imported if it passes the validation checks.

() NOTE: Although you can export the factory-installed license, you cannot import it. To import the license, download the equivalent license from the Digital Locker or retrieve it from the email you received when you purchased the license.

(i) **NOTE:** After importing the license, you need to re-login to iDRAC. This is applicable only for iDRAC web interface.

- Export Exports the installed license. For more information, see the *iDRAC Online Help*.
- Delete Deletes the license. For more information, see the *iDRAC Online Help*.
- Learn More Learn more about an installed license, or the licenses available for a component installed in the server.

NOTE: For the Learn More option to display the correct page, ensure that ***.dell.com** is added to the list of Trusted Sites in the Security Settings. For more information, see the Internet Explorer help documentation.

For one-to-many license deployment, you can use Dell License Manager. For more information, see *Dell License Manager User's Guide* available at https://www.dell.com/esmmanuals.

Following are the user privilege requirements for different license operation:

- Licence View and Export: Login privilege.
- License Import and Delete: Login + Configure iDRAC + Server Control privilege.

Managing licenses using iDRAC web interface

To manage the licenses using the iDRAC web interface, go to **Configuration** > **Licenses**.

The **Licensing** page displays the licenses that are associated to devices, or the licenses that are installed but the device is not present in the system. For more information on importing, exporting, or deleting a license, see the *iDRAC Online Help*.

Managing licenses using RACADM

To manage licenses using RACADM, use the license subcommand. For more information, see the

iDRAC RACADM CLI Guide available at https://www.dell.com/idracmanuals.

Fls. <u>411</u> Mov. 32



Licensed features in iDRAC9

The following table lists iDRAC9 features that are enabled based on the license purchased:

Table 2. Licensed features in iDRAC9

Feature	iDRAC 9 Basic	iDRAC9 Express	iDRAC9 Express for Blades	iDRAC9 Enterprise	iDRAC9 Datacenter
Interfaces / Standards	-			•	`
iDRAC RESTful API and Redfish	Yes	Yes	Yes	Yes	Yes
IPMI 2.0	Yes	Yes	Yes	Yes	Yes
DCMI 1.5	Yes	Yes	Yes	Yes	Yes
Web-based GUI	Yes	Yes	Yes	Yes	Yes
RACADM command line (local/remote)	Yes	Yes	Yes	Yes	Yes
SSH	Yes	Yes	Yes	Yes	Yes
Serial Redirection	Yes	Yes	Yes	Yes	Yes
WSMan	Yes	Yes	Yes	Yes	Yes
Network Time Protocol	No	Yes	Yes	Yes	Yes
Connectivity		-	-		
Shared NIC (LOM)	Yes	Yes	N/A	Yes	Yes
Dedicated NIC	Yes	Yes	Yes	Yes	Yes
VLAN tagging	Yes	Yes	Yes	Yes	Yes
IPv4	Yes	Yes	Yes	Yes	Yes
IPv6	Yes	Yes	Yes	Yes	Yes
DHCP	Yes	Yes	Yes	Yes	Yes
DHCP with zero touch	No	No	No	Yes	Yes
Dynamic DNS	Yes	Yes	Yes	Yes	Yes
OS pass-through	Yes	Yes	Yes	Yes	Yes
iDRAC Direct -Front panel USB	Yes	Yes	Yes	Yes	Yes
Connection View	Yes	Yes	No	Yes	Yes
Security		-	-		
Role-based authority	Yes	Yes	Yes	Yes	Yes
Local users	Yes	Yes	Yes	Yes	Yes
SSL encryption	Yes	Yes	Yes	Yes	Yes
Secure Enterprise Key Manager	No	No	No	Yes (with SEKM license)	Yes (with SEKM license)
IP blocking	No	Yes	Yes	Yes	Yes
Directory services (AD, LDAP)	No	No	No	Yes	Yes



Table 2. Licensed features in iDRAC9 (continued)

Feature	iDRAC 9 Basic	iDRAC9 Express	iDRAC9 Express for Blades	iDRAC9 Enterprise	iDRAC9 Datacenter
Two-factor authentication (smart card)	No	No	No	Yes	Yes
Single sign-On	No	No	No	Yes	Yes
PK authentication (for SSH)	No	Yes	Yes	Yes	Yes
OAuth integration with Web based Authentication services	No	No	No	No	Yes
OpenID Connect for Dell EMC Consoles	No	No	No	No	Yes
FIPS 140-2	Yes	Yes	Yes	Yes	Yes
Secure UEFI boot - certificate management	Yes	Yes	Yes	Yes	Yes
Lock down mode	No	No	No	Yes	Yes
Unique iDRAC default password	Yes	Yes	Yes	Yes	Yes
Customizable Security Policy Banner - login page	Yes	Yes	Yes	Yes	Yes
Easy Multi Factor Authentication	No	No	No	No	Yes
Auto Certificate Enrollment (SSL Certs)	No	No	No	No	Yes
iDRAC Quick Sync 2 - optional auth for read operations	Yes	Yes	Yes	Yes	Yes
iDRAC Quick Sync 2 - add mobile device number to LCL	Yes	Yes	Yes	Yes	Yes
System Erase of internal storage devices	Yes	Yes	Yes	Yes	Yes
Remote Presence					
Power control	Yes	Yes	Yes	Yes	Yes
Boot control	Yes	Yes	Yes	Yes	Yes
Serial-over-LAN	Yes	Yes	Yes	Yes	Yes
Virtual Media	No	No	Yes	Yes	Yes
Virtual Folders	No	No	No	Yes	Yes
Remote File Share	No	No	No	Yes	Yes
HTML5 access to Virtual Console	No	No	Yes	Yes	Yes
Virtual Console	No	No	Yes	Yes	Yes
VNC connection to OS	No	No	No	Yes	Yes
Quality/bandwidth control	No	No	No	Yes	Yes



Table 2. Licensed features in iDRAC9 (continued)

Feature	iDRAC 9 Basic	iDRAC9 Express	iDRAC9 Express for Blades	iDRAC9 Enterprise	iDRAC9 Datacenter
Virtual Console collaboration (up to six simultaneous users)	No	No	No (One user only)	Yes	Yes
Virtual Console chat	No	No	No	Yes	Yes
Virtual Flash partitions	No	No	No	Yes	Yes
(i) NOTE: vFlash is not av	/ailable in i	iDRAC9 for	PowerEdge Rx5xx/Cx	ōxx.	
Group Manager	No	No	No	Yes	Yes
HTTP / HTTPS support along with NFS/CIFS	Yes	Yes	Yes	Yes	Yes
Power and Thermal	<u>. </u>			1	
Real-time power meter	Yes	Yes	Yes	Yes	Yes
Power thresholds and alerts	No	Yes	Yes	Yes	Yes
Real-time power graphing	No	Yes	Yes	Yes	Yes
Historical power counters	No	Yes	Yes	Yes	Yes
Power capping	No	No	No	Yes	Yes
Power Center integration	No	No	No	Yes	Yes
Temperature monitoring	Yes	Yes	Yes	Yes	Yes
Temperature graphing	No	Yes	Yes	Yes	Yes
PCIe airflow customization (LFM)	No	No	No	No	Yes
Custom Exhaust Control	No	No	No	No	Yes
Custom Delta-T control	No	No	No	No	Yes
System Airflow Consumption	No	No	No	No	Yes
Custom PCle inlet temperature	No	No	No	No	Yes
Health Monitoring					
Full agent-free monitoring	Yes	Yes	Yes	Yes	Yes
Predictive failure monitoring	Yes	Yes	Yes	Yes	Yes
SNMPv1, v2, and v3 (traps and gets)	Yes	Yes	Yes	Yes	Yes
Email Alerting	No	Yes	Yes	Yes	Yes
Configurable thresholds	Yes	Yes	Yes	Yes	Yes
Fan monitoring	Yes	Yes	Yes	Yes	Yes

28 Overview of iDRAC


Table 2. Licensed features in iDRAC9 (continued)

Feature	iDRAC 9 Basic	iDRAC9 Express	iDRAC9 Express for Blades	iDRAC9 Enterprise	iDRAC9 Datacenter
Power Supply monitoring	Yes	Yes	Yes	Yes	Yes
Memory monitoring	Yes	Yes	Yes	Yes Yes	
CPU monitoring	Yes	Yes	Yes	Yes	Yes
RAID monitoring	Yes	Yes	Yes	Yes	Yes
NIC monitoring	Yes	Yes	Yes	Yes	Yes
Optic Inventory	Yes	Yes	Yes	Yes	Yes
Optic Statistics	No	No	No	No	Yes
HD monitoring (enclosure)	Yes	Yes	Yes	Yes	Yes
Out of Band Performance Monitoring	No	No	No	Yes	Yes
Alerts for excessive SSD wear	Yes	Yes	Yes	Yes	Yes
Customizable settings for Exhaust Temperature	Yes	Yes	Yes	Yes Yes	
Serial Console Logs	No	No	No	No	Yes
SMART logs for Storage Drives	No	No	No	No	Yes
Idle Server detection	No	No	No	No	Yes
Telemetry Streaming	No	No	No	No	Yes
(i) NOTE: The OpenMana iDRAC.	ige Enterp	rise Advanc	ed license and the Pow	verManage Plugin support telemet	ry data pulls from the
Update					
Remote agent-free update	Yes	Yes	Yes	Yes	Yes
Embedded update tools	Yes	Yes	Yes	Yes	Yes
Update from repository (Auto-Update)	No	No	No	Yes	Yes
Schedule update from repository	No	No	No	Yes	Yes
Improved PSU firmware updates	Yes	Yes	Yes	Yes	Yes
Deployment and Configu	iration				
Local configuration via F10	Yes	Yes	Yes	Yes	Yes



Table 2. Licensed features in iDRAC9 (continued)

Feature	iDRAC 9 Basic	iDRAC9 Express	iDRAC9 Express for Blades	iDRAC9 Enterprise	iDRAC9 Datacenter
Embedded OS deployment tools	Yes	Yes	Yes	Yes	Yes
Embedded configuration tools	Yes	Yes	Yes	Yes	Yes
Auto-Discovery	No	Yes	Yes	Yes	Yes
Remote OS deployment	No	Yes	Yes	Yes	Yes
Embedded driver pack	Yes	Yes	Yes	Yes	Yes
Full configuration inventory	Yes	Yes	Yes	Yes	Yes
Inventory export	Yes	Yes	Yes	Yes	Yes
Remote configuration	Yes	Yes	Yes	Yes	Yes
Zero-touch configuration	No	No	No	Yes	Yes
System Retire/Repurpose	Yes	Yes	Yes	Yes	Yes
Server Configuration Profile in GUI	Yes	Yes	Yes	Yes Yes	
Add BIOS configuration to iDRAC GUI	Yes	Yes	Yes	Yes Yes	
GPU properties	No	No	No	Yes	Yes
Diagnostics, Service, and	d Logging]			
Embedded diagnostic tools	Yes	Yes	Yes	Yes	Yes
Part Replacement	No	Yes	Yes	Yes	Yes
NOTE: After performin configuration, Lifecycle	ig part rep Logs rep	lacement of orts double	n RAID hardware, and t part replacement entri	the process is complete for replac es which is expected behavior.	ing firmware and
Easy Restore (system configuration)	Yes	Yes	Yes	Yes	Yes
Easy Restore Auto Timeout	Yes	Yes	Yes	Yes	Yes
i NOTE: Server Backup	and Restc	ore features	are not available in iDR	AC9 for PowerEdge Rx5xx/Cx5x	х.
LED Health status indicators	Yes	Yes	N/A	Yes	Yes
LCD screen (iDRAC9 requires optional)	Yes	Yes	N/A	Yes	Yes
iDRAC Quick Sync 2 (BLE/Wi-Fi hardware)	Yes	Yes	Yes	Yes	Yes

Inserido ao protocolo **22.951.206-4** por: **Pedro Henrique Golin Linhares** em: 31/03/2025 12:49. A autenticidade deste documento pode ser validada no endereço: https://www.eprotocolo.pr.gov.br/spiweb/validarDocumento com o código: **692a5ef82beb173ed33dad4ee6505546**.



Table 2. Licensed features in iDRAC9 (continued)

Feature	iDRAC 9 Basic	iDRAC9 Express	iDRAC9 Express for Blades	iDRAC9 Enterprise	iDRAC9 Datacenter
iDRAC Direct (front USB management port)	Yes	Yes	Yes	Yes	Yes
iDRAC Service Module (iSM) embedded	Yes	Yes	Yes	Yes Yes	
iSM to in-band alert forwarding to consoles	Yes	Yes	Yes	Yes	Yes
SupportAssist Collection (embedded)	Yes	Yes	Yes	Yes	Yes
Crash screen capture	No	Yes	Yes	Yes	Yes
Crash video capture ¹	No	No	No	Yes	Yes
Agent Free Crash Video Capture (Windows only)	No	No	No	No	Yes
Boot capture	No	No	No	Yes	Yes
Manual reset for iDRAC (LCD ID button)	Yes	Yes	Yes	Yes	Yes
Remote reset for iDRAC (requires iSM)	Yes	Yes	Yes	Yes Yes	
Virtual NMI	Yes	Yes	Yes	Yes	Yes
OS watchdog	Yes	Yes	Yes	Yes	Yes
System Event Log	Yes	Yes	Yes	Yes	Yes
Lifecycle Log	Yes	Yes	Yes	Yes	Yes
Enhanced Logging in Lifecycle Controller Log	Yes	Yes	Yes	Yes	Yes
Work notes	Yes	Yes	Yes	Yes	Yes
Remote Syslog	No	No	No	Yes	Yes
License management	Yes	Yes	Yes	Yes	Yes
Improved Customer Exp	erience		•	•	•
iDRAC -Faster processor, more memory	N/A	Yes	N/A	Yes	Yes
GUI rendered in HTML5	N/A	Yes	N/A	Yes	Yes
Add BIOS configuration to iDRAC GUI	N/A	Yes	N/A	Yes	Yes

[1] Requires iSM or OMSA agent on target server.

Interfaces and protocols to access iDRAC

The following table lists the interfaces to access iDRAC.



(i) NOTE: Using more than one interface at the same time may generate unexpected results.

Table 3. Interfaces and protocols to access iDRAC

Interface or Protocol	Description					
iDRAC Settings Utility (F2)	Use the iDRAC Settings utility to perform pre-OS operations. It has a subset of the features that are available in iDRAC web interface along with other features.					
	To access iDRAC Settings utility, press <f2> during boot and then click iDRAC Settings on the System Setup Main Menu page.</f2>					
Lifecycle Controller (F10)	Use Lifecycle Controller to perform iDRAC configurations. To access Lifecycle Controller, press <f10> during boot and go to System Setup > Advanced Hardware Configuration > iDRAC Settings. For more information, see <i>Lifecycle Controller User's Guide</i> available at dell.com/ idracmanuals.</f10>					
iDRAC Web Interface	Use the iDRAC web interface to manage iDRAC and monitor the managed system. The browser connects to the web server through the HTTPS port. Data streams are encrypted using 128-bit SSL to provide privacy and integrity. Any connection to the HTTP port is redirected to HTTPS. Administrators can upload their own SSL certificate through an SSL CSR generation process to secure the web server. The default HTTP and HTTPS ports can be changed. The user access is based on user privileges.					
OpenManage Enterprise (OME)	(i) NOTE: This interface is only available for MX platforms.					
Enterprise (OME) Modular Web Interface	 In addition to monitoring and managing the chassis, use the OME-Modular web interface to: View the status of a managed system Update iDRAC firmware Configure iDRAC network settings Log in to iDRAC web interface 					
	Start, stop, or reset the managed system					
	For more information, see the OME - Modular for PowerEdge MX7000 Chassis User's Guide available at https://www.dell.com/openmanagemanuals.					
CMC Web Interface	() NOTE: This interface is not available in MX platforms.					
	 In addition to monitoring and managing the chassis, use the CMC web interface to: View the status of a managed system Update iDRAC firmware Configure iDRAC network settings Log in to iDRAC web interface Start, stop, or reset the managed system Update BIOS, PERC, and supported network adapters 					
Server LCD Panel/ Chassis LCD Panel	 Use the LCD on the server front panel to: View alerts, iDRAC IP or MAC address, user programmable strings. Set DHCP Configure iDRAC static IP settings. 					
	For blade servers, the LCD is on the chassis front panel and is shared between all the blades.					
	To reset iDRAC without rebooting the server, press and hold the System Identification button $oldsymbol{\Theta}$ for 16 seconds.					
	() NOTE: LCD panel is only available with rack or tower systems that support front bezel. For blade servers, the LCD is on the chassis front panel and is shared between all the blades.					
RACADM	 Use this command-line utility to perform iDRAC and server management. You can use RACADM locally and remotely. Local RACADM command-line interface runs on the managed systems that have Server Administrator installed. Local RACADM communicates with iDRAC through its in-band IPMI host 					

Inserido ao protocolo **22.951.206-4** por: **Pedro Henrique Golin Linhares** em: 31/03/2025 12:49. A autenticidade deste documento pode ser validada no endereço: https://www.eprotocolo.pr.gov.br/spiweb/validarDocumento com o código: **692a5ef82beb173ed33dad4ee6505546**.



Table 3. Interfaces and protocols to access iDRAC (continued)

Interface or Protocol	Description				
	 interface. Since it is installed on the local managed system, users are required to log in to the operating system to run this utility. A user must have a full administrator privilege or be a root user to use this utility. Remote RACADM is a client utility that runs on a management station. It uses the out-of-band network interface to run RACADM commands on the managed system and uses the HTTPs channel. The -r option runs the RACADM command over a network. Firmware RACADM is accessible by logging in to iDRAC using SSH. You can run the firmware RACADM commands without specifying the iDRAC IP, user name, or password. You do not have to specify the iDRAC IP, user name, or password to run the firmware RACADM commands. After you enter the RACADM prompt, you can directly run the commands without the racadm prefix. 				
iDRAC RESTful API and Redfish	The Redfish Scalable Platforms Management API is a standard defined by the Distributed Management Task Force (DMTF). Redfish is a next-generation systems management interface standard, which enables scalable, secure, and open server management. It is a new interface that uses RESTful interface semantics to access data that is defined in model format to perform out- of-band systems management. It is suitable for a wide range of servers ranging from stand-alone servers to rack mount and bladed environments and for large scale cloud environments. Redfish provides the following benefits over existing server management methods: Increased simplicity and usability High data security Programmable interface that can be easily scripted Following widely used standards				
	For iDRAC Redfish API guide, go to www.api-marketplace.com				
WSMan	The LC-Remote Service is based on the WSMan protocol to do one-to-many systems management tasks. You must use WSMan client such as WinRM client (Windows) or the OpenWSMan client (Linux) to use the LC-Remote Services functionality. You can also use Power Shell or Python to script to the WSMan interface.				
	Web Services for Management (WSMan) is a Simple Object Access Protocol (SOAP)-based protocol used for systems management. iDRAC uses WSMan to convey Distributed Management Task Force (DMTF) Common Information Model (CIM)-based management information. The CIM information defines the semantics and information types that can be modified in a managed system. The data available through WSMan is provided by iDRAC instrumentation interface mapped to the DMTF profiles and extension profiles.				
	 For more information, see the following: Lifecycle Controller Remote Services Quick Start Guide available at https://www.dell.com/ idracmanuals. MOFs and Profiles — http://downloads.dell.com/wsman. DMTF website — dmtf.org/standards/profiles 				
SSH	Use SSH to run RACADM commands. The SSH service is enabled by default on iDRAC. The SSH service can be disabled in iDRAC. iDRAC only supports SSH version 2 with the RSA host key algorithm. A unique 1024-bit RSA host key is generated when you power-up iDRAC for the first time.				
IPMITool	Use the IPMITool to access the remote system's basic management features through iDRAC. The interface includes local IPMI, IPMI over LAN, IPMI over Serial, and Serial over LAN. For more information on IPMITool, see the <i>Dell OpenManage Baseboard Management Controller Utilities User's Guide</i> at dell.com/idracmanuals.				
NTLM	iDRAC allows NTLM to provide authentication, integrity, and confidentiality to the users. NT LAN Manager (NTLM) is a suite of Microsoft security protocols and it works in a Windows network.				
SMB	iDRAC9 supports the Server Message Block (SMB) Protocol. This is a network file sharing protocol and the default minimum SMB version supported is 2.0, SMBv1 is no longer supported.				



Table 3. Interfaces and protocols to access iDRAC (continued)

Interface or Protocol	Description
NFS	iDRAC9 supports Network File System (NFS). This is a distributed filesystem protocol that enables users to mount remote directories on the servers.

iDRAC port information

The following table lists the ports that are required to remotely access iDRAC through firewall. These are the default ports iDRAC listens to for connections. Optionally, you can modify most of the ports. To modify ports, see Configuring services.

Table 4. Port	s iDRAC listens	s for connections
---------------	-----------------	-------------------

Port number	Туре	Function	Configurabl e port	Maximum Encryption Level	
22	TCP	SSH	Yes	256-bit SSL	
80	TCP	HTTP	Yes	None	
161	UDP	SNMP Agent	Yes	None	
443	TCP	 Web GUI access with HTTPS Virtual Console and Virtual Media with eHTML5 option Virtual Console and Virtual Media with HTML5 option when web server redirection is enabled 	Yes	256-bit SSL	
623	UDP	RMCP/RMCP+	No	128-bit SSL	
5000	TCP	iDRAC to iSM	No	256-bit SSL	
(i) NOTE: Maximum encryption level is 256-bit SSL if both iSM 3.4 or higher and iDRAC firmware 3.30.30.30 or higher are installed.					
5900	ТСР	Virtual console and virtual media with HTML5, Java and ActiveX option	Yes	128-bit SSL	
5901	TCP	VNC	Yes	128-bit SSL	
	Port 5901	opens when VNC feature is enabled.			

The following table lists the ports that iDRAC uses as a client:

Table 5. Ports iDRAC uses as client

Port number	Туре	Function	Configurable port	Maximum Encryption Level
25	TCP	SMTP	Yes	None
53	UDP	DNS	No	None
68	UDP	DHCP-assigned IP address	No	None
69	TFTP	TFTP	No	None
123	UDP	Network Time Protocol (NTP)	No	None
162	UDP	SNMP trap	Yes	None
445	TCP	Common Internet File System (CIFS)	No	None
636	TCP	LDAP Over SSL (LDAPS)	No	256-bit SSL
2049	TCP	Network File System (NFS)	No	None



Table 5. Ports iDRAC uses as client (continued)

Port number	Туре	Function	Configurable port	Maximum Encryption Level		
3269	TCP	LDAPS for global catalog (GC)	No	256-bit SSL		
5353	UDP	mDNS	No	None		
(j) NOTE: V 5353. Ho in the po	() NOTE: When node initiated discovery or Group Manager is enabled, iDRAC uses mDNS to communicate through port 5353. However, when both are disabled, port 5353 is blocked by iDRAC's internal firewall and appears as open filtered port in the port scans.					
514	UDP	Remote syslog	Yes	None		

Other documents you may need

Some of the iDRAC interfaces have the integrated *Online Help* document that can be accessed by clicking on the help (?) icon. The *Online Help* provides detailed information about the fields available on the web interface and the descriptions for the same. In addition, following documents are available on the Dell Support website at **dell.com/support** that provide additional information about the setup and operation of iDRAC in your system.

- The iDRAC Redfish API Guide available at https://developer.dell.com provides information about Redfish API.
- The *iDRAC RACADM CLI Guide* provides information about the RACADM sub-commands, supported interfaces, and iDRAC property database groups and object definitions.
- The Systems Management Overview Guide provides brief information about the various software available to perform systems management tasks.
- The *Dell Remote Access Configuration Tool User's Guide* provides information on how to use the tool to discover iDRAC IP addresses in your network and perform one-to-many firmware updates and active directory configurations for the discovered IP addresses.
- The *Dell Systems Software Support Matrix* provides information about the various Dell systems, the operating systems supported by these systems, and the Dell OpenManage components that can be installed on these systems.
- The iDRAC Service Module User's Guide provides information to install the iDRAC Service Module.
- The Dell OpenManage Server Administrator Installation Guide contains instructions to help you install Dell OpenManage Server Administrator.
- The Dell OpenManage Management Station Software Installation Guide contains instructions to help you install Dell OpenManage management station software that includes Baseboard Management Utility, DRAC Tools, and Active Directory Snap-In.
- The Dell OpenManage Baseboard Management Controller Management Utilities User's Guide has information about the IPMI interface.
- The *Release Notes* provides last-minute updates to the system or documentation or advanced technical reference material intended for experienced users or technicians.

The following system documents are available to provide more information:

- The safety instructions that came with your system provide important safety and regulatory information. For additional regulatory information, see the Regulatory Compliance home page at **dell.com/regulatory_compliance**. Warranty information may be included within this document or as a separate document.
- The Rack Installation Instructions included with your rack solution describe how to install your system into a rack.
- The *Getting Started Guide* provides an overview of system features, setting up your system, and technical specifications.
- The *Installation and Service Manual* provides information about system features and describes how to troubleshoot the system and install or replace system components.

Contacting Dell

NOTE: If you do not have an active Internet connection, you can find contact information on your purchase invoice, packing slip, bill, or Dell product catalog.

Dell provides several online and telephone-based support and service options. Availability varies by country and product, and some services may not be available in your area. To contact Dell for sales, technical support, or customer service issues, visit https://www.dell.com/contactdell.



Accessing documents from Dell support site

You can access the required documents in one of the following ways:

- Using the following links:
 - For all Enterprise Systems Management and OpenManage Connections documents https://www.dell.com/ esmmanuals
 - For OpenManage documents https://www.dell.com/openmanagemanuals
 - For iDRAC and Lifecycle Controller documents https://www.dell.com/idracmanuals
 - For Serviceability Tools documents https://www.dell.com/serviceabilitytools
 - For Client Command Suite Systems Management documents https://www.dell.com/omconnectionsclient

Accessing documents using the product search

- 1. Go to https://www.dell.com/support.
- 2. In the Enter a Service Tag, Serial Number... search box, type the product name. For example, PowerEdge or iDRAC.

A list of matching products is displayed.

- 3. Select your product and click the search icon or press enter.
- 4. Click **DOCUMENTATION**.
- 5. Click MANUALS AND DOCUMENTS.

Accessing documents using product selector

You can also access documents by selecting your product.

- 1. Go to https://www.dell.com/support.
- 2. Click Browse all products.
- 3. Click the desired product category, such as Servers, Software, Storage, and so on.
- 4. Click the desired product and then click the desired version if applicable.

(i) NOTE: For some products, you may need to navigate through the subcategories.

- 5. Click **DOCUMENTATION**.
- 6. Click MANUALS AND DOCUMENTS.

Accessing Redfish API Guide

The Redfish API guide is now available at the Dell API Marketplace. To access the Redfish API guide:

- 1. Go to www.api-marketplace.com.
- 2. Click Explore API and then click APIs.
- 3. Under iDRAC9 Redfish API, click View More.



Logging in to iDRAC

You can log in to iDRAC as an iDRAC user, a Microsoft Active Directory user, or a Lightweight Directory Access Protocol (LDAP) user. You can also log in using OpenID Connect and Single Sign-On or Smart Card.

To improve security, each system is shipped with a unique password for iDRAC, which is available on the system information tag. This unique password improves security of iDRAC and your server. The default user name is *root*.

While ordering the system, you can choose to retain the legacy password—calvin—as the default password. If you choose to retain the legacy password, the password is not available on the system information tag.

In this version, DHCP is enabled by default and iDRAC IP address is assigned dynamically.

() NOTE:

- You must have login to iDRAC privilege to log in to iDRAC.
- iDRAC GUI does not support browser buttons such as **Back**, **Forward**, or **Refresh**.

NOTE: For information about recommended characters for user names and passwords, see Recommended characters in user names and passwords.

To change the default password, see Changing the default login password.

Customizable security banner

You can customize the security notice that is displayed on the login page. You can use SSH, RACADM, Redfish, or WSMan to customize the notice. Depending on the language you use, the notice can be either 1024 or 512 UTF-8 characters long.

OpenID Connect

(i) NOTE: This feature is only available for MX platforms.

You can log in to iDRAC using credentials of other web consoles such as Dell EMC OpenManage Enterprise (OME) - Modular. When this feature is enabled, the console starts managing the user permissions on the iDRAC. iDRAC provides the user session with all the permissions that are specified by the console.

(i) NOTE: When lockdown mode is enabled, OpenID Connect login options are not displayed in iDRAC login page.

You can now get access to detailed help without logging in to iDRAC. Use the links on the iDRAC login page to access help and version information, drivers and downloads, manuals and TechCenter.

Topics:

- Force Change of Password (FCP)
- Logging into iDRAC using OpenID Connect
- Logging in to iDRAC as local user, Active Directory user, or LDAP user
- Logging in to iDRAC as a local user using a smart card
- Logging in to iDRAC using Single Sign-On
- Accessing iDRAC using remote RACADM
- Accessing iDRAC using local RACADM
- Accessing iDRAC using firmware RACADM
- Simple 2-Factor Authentication (Simple 2FA)
- RSA SecurID 2FA
- Viewing system health
- Logging in to iDRAC using public key authentication
- Multiple iDRAC sessions

Logging in to iDRAC 37



- Secure default password
- Changing the default login password
- Enabling or disabling default password warning message
- Password Strength Policy
- IP Blocking
- Enabling or disabling OS to iDRAC Pass-through using web interface
- Enabling or disabling alerts using RACADM

Force Change of Password (FCP)

The 'Force Change of Password' feature prompts you to change the factory default password of the device. The feature can be enabled as part of factory configuration.

The FCP screen appears after successful user authentication and cannot be skipped. Only after the user enters a password, normal access and operation will be allowed. The state of this attribute will not be affected by a 'Reset Configuration to Defaults' operation.

(i) NOTE: To set or reset the FCP attribute, you must have Login privilege and User configuration privilege.

(i) NOTE: When FCP is enabled, 'Default Password Warning' setting is disabled after changing the default user password.

(i) NOTE: When root user logs in via Public Key Authentication (PKA), FCP is bypassed.

When FCP is enabled, following actions are not allowed:

- Login to iDRAC through any UI except IPMIpover-LAN interface which uses CLI with default user credentials.
- Login to iDRAC through OMM app via Quick Sync-2
- Add a member iDRAC in Group Manager.

Logging into iDRAC using OpenID Connect

(i) NOTE: This feature is only available in MX platforms.

To log in to iDRAC using the OpenID Connect:

- In a supported web browser, type https://[iDRAC-IP-address] and press Enter. The Login page is displayed.
- 2. Select OME Modular from the Log In with: menu. The console login page is displayed.
- 3. Enter the console User name and Password.
- 4. Click Log in.

You are logged in to iDRAC with the console user privileges.

(i) NOTE: When lockdown mode is enabled, OpenID Connect login option is not displayed in iDRAC login page.

Logging in to iDRAC as local user, Active Directory user, or LDAP user

Before you log in to iDRAC using the web interface, ensure that you have configured a supported web browser and the user account is created with the required privileges.

(i) NOTE: The user name is not case-sensitive for an Active Directory user. The password is case-sensitive for all users.

(i) NOTE: In addition to Active Directory, openLDAP, openDS, Novell eDir, and Fedora-based directory services are supported.

(i) NOTE: LDAP authentication with OpenDS is supported. The DH key must be larger than 768 bits.



NOTE: RSA feature can be configured and enabled for LDAP user, but the RSA does not support if the LDAP is configured on Microsoft active directory. Hence LDAP user login fails. RSA is supported only for OpenLDAP.

To log in to iDRAC as local user, Active Directory user, or LDAP user:

- 1. Open a supported web browser.
- 2. In the Address field, type https://[iDRAC-IP-address] and press Enter.
 - () NOTE: If the default HTTPS port number (port 443) changes, enter: https://[iDRAC-IP-address]:[port-number] where [iDRAC-IP-address] is the iDRAC IPv4 or IPv6 address and [port-number] is the HTTPS port number.

The **Login** page is displayed.

- 3. For a local user:
 - In the **Username** and **Password** fields, enter your iDRAC user name and password.
 - From the **Domain** drop-down menu, select **This iDRAC**.
- 4. For an Active Directory user, in the User name and Password fields, enter the Active Directory user name and password. If you have specified the domain name as a part of the username, select This iDRAC from the drop-down menu. The format of the user name can be: <domain>\<username>, <domain>\<username>, or <user>@<domain>.

For example, dell.com\john_doe, or JOHN_DOE@DELL.COM.

If the domain is not specified in the user name, select the Active Directory domain from the **Domain** drop-down menu.

- 5. For an LDAP user, in the **Username** and **Password** fields, enter your LDAP user name and password. Domain name is not required for LDAP login. By default, **This iDRAC** is selected in the drop-down menu.
- 6. Click **Submit**. You are logged in to iDRAC with the required user privileges.

If you log in with Configure Users privileges and the default account credentials, and if the default password warning feature is enabled, the **Default Password Warning** page is displayed allowing you to easily change the password.

Logging in to iDRAC as a local user using a smart card

Before you log in as a local user using Smart Card, make sure to:

- Upload user smart card certificate and the trusted Certificate Authority (CA) certificate to iDRAC.
- Enable smart card logon.

The iDRAC web interface displays the smart card logon page for users who are configured to use the smart card.

- **NOTE:** Depending on the browser settings, you are prompted to download and install the smart card reader ActiveX plug-in when using this feature for the first time.
- To log in to iDRAC as a local user using a smart card:
- 1. Access the iDRAC web interface using the link https://[IP address].

The **iDRAC Login** page is displayed prompting you to insert the smart card.

- **NOTE:** If the default HTTPS port number (port 443) changes, type: https://[IP address]:[port number] where, [IP address] is the IP address for the iDRAC and [port number] is the HTTPS port number.
- 2. Insert the smart card into the reader and click **Login**.
 - A prompt is displayed for the smart card's PIN. A password is not required.
- 3. Enter the Smart Card PIN for local smart card users.
 - You are logged in to the iDRAC.
 - (i) **NOTE:** If you are a local user for whom **Enable CRL check for Smart Card Logon** is enabled, iDRAC attempts to download the certificate revocation list (CRL) and checks the CRL for the user's certificate. The login fails if the certificate is listed as revoked in the CRL or if the CRL cannot be downloaded for some reason.

(i) NOTE: If you log in to iDRAC using smart card when RSA is enabled, RSA token is bypassed and you can login directly.



Logging in to iDRAC as an Active Directory user using a smart card

Before you log in as an Active Directory user using smart card, ensure that you:

- Upload a Trusted Certificate Authority (CA) certificate (CA-signed Active Directory certificate) to iDRAC.
- Configure the DNS server.
- Enable Active Directory login.
- Enable smart card login.

To log in to iDRAC as an Active Directory user using smart card:

1. Log in to iDRAC using the link https://[IP address].

The **iDRAC Login** page is displayed prompting you to insert the smart card.

NOTE: If the default HTTPS port number (port 443) is changed, type: https://[IP address]:[port number] where, [IP address] is the iDRAC IP address and [port number] is the HTTPS port number.

- 2. Insert the smart card and click Login.
 - A prompt is displayed for the smart card's **PIN**.
- **3.** Enter the PIN and click **Submit**.

You are logged in to iDRAC with your Active Directory credentials.

(i) NOTE:

If the smart card user is present in Active Directory, an Active Directory password is not required.

Logging in to iDRAC using Single Sign-On

When Single Sign-On (SSO) is enabled, you can log in to iDRAC without entering your domain user authentication credentials, such as user name and password.

(i) NOTE: When AD user configures SSO while RSA is enabled, the RSA token is bypassed and user logs in directly.

Logging in to iDRAC SSO using iDRAC web interface

Before logging in to iDRAC using Single Sign-On, ensure that:

- You have logged in to your system using a valid Active Directory user account.
- Single Sign-On option is enabled during Active Directory configuration.
- To log in to iDRAC using web interface:
- 1. Log in to your management station using a valid Active Directory account.
- 2. In a web browser, type https://[FQDN address].
 - () NOTE: If the default HTTPS port number (port 443) has been changed, type: https://[FQDN address]:[port number] where [FQDN address] is the iDRAC FQDN (iDRACdnsname.domain. name) and [port number] is the HTTPS port number.
 - (i) NOTE: If you use IP address instead of FQDN, SSO fails.

iDRAC logs you in with appropriate Microsoft Active Directory privileges, using your credentials that were cached in the operating system when you logged in using a valid Active Directory account.

Logging in to iDRAC SSO using CMC web interface

(i) NOTE: This feature is not available on MX platforms.

Using the SSO feature, you can launch the iDRAC web interface from the CMC web interface. A CMC user has the CMC user privileges when launching iDRAC from CMC. If the user account is present in CMC and not in iDRAC, the user can still launch iDRAC from CMC.

If iDRAC network LAN is disabled (LAN Enabled = No), SSO is not available.

40 Logging in to iDRAC



If the server is removed from the chassis, iDRAC IP address is changed, or there is a problem in iDRAC network connection, the option to Launch iDRAC is grayed-out in the CMC web interface.

For more information, see the Chassis Management Controller User's Guide available at https://www.dell.com/cmcmanuals.

Accessing iDRAC using remote RACADM

You can use remote RACADM to access iDRAC using RACADM utility.

For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

If the management station has not stored the iDRAC's SSL certificate in its default certificate storage, a warning message is displayed when you run the RACADM command. However, the command is executed successfully.

() NOTE: The iDRAC certificate is the certificate iDRAC sends to the RACADM client to establish the secure session. This certificate is either issued by a CA or self-signed. In either case, if the management station does not recognize the CA or signing authority, a warning is displayed.

Validating CA certificate to use remote RACADM on Linux

Before running remote RACADM commands, validate the CA certificate that is used for secure communications.

To validate the certificate for using remote RACADM:

1. Convert the certificate in DER format to PEM format (using openssl command-line tool):

openssl x509 -inform pem -in [yourdownloadedderformatcert.crt] -outform pem -out [outcertfileinpemformat.pem] -text

- Find the location of the default CA certificate bundle on the management station. For example, for RHEL5 64 bit, it is /etc/pki/tls/cert.pem.
- 3. Append the PEM formatted CA certificate to the management station CA certificate. For example, use the cat command: cat testcacert.pem >> cert.pem
- **4.** Generate and upload the server certificate to iDRAC.

Accessing iDRAC using local RACADM

For information to access iDRAC using local RACADM, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/ idracmanuals.

Accessing iDRAC using firmware RACADM

You can use SSH interface to access iDRAC and run firmware RACADM commands. For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Simple 2-Factor Authentication (Simple 2FA)

iDRAC offers simple 2-factor authentication option to enhance the security to the local users for logging in. When you log in from a source IP-address which is different from the last login, you will be prompted to enter the second factor authentication details.

Simple two factor authentication has two steps of authentication:

- iDRAC User name and password
- Simple 6 digit code which is sent to the user via email. User needs to enter this 6 digit code when prompted at login.

() NOTE:

• To receive 6 digit code, it is mandatory to configure 'Custom Sender Address' and have valid SMTP configuration.



- The 2FA code expires after 10 minutes or is invalidated if it is already consumed before expiry.
- If a user attempts to login from another location with a different IP-Address while a pending 2FA challenge for the original IP-Address is still outstanding, the same token will be sent for login attempt from the new IP address.
- The feature is supported with iDRAC Enterprise or Datacenter license.

When 2FA is enabled, following actions are not allowed:

- Login to iDRAC through any UI which uses CLI with default user credentials.
- Login to iDRAC through OMM app via Quick Sync-2
- Add a member iDRAC in Group Manager.

NOTE: Racadm, Redfish, WSMAN, IPMI LAN, Serial, CLI from a source IP works only after successful login from supported interfaces like iDRAC GUI and SSH.

RSA SecurID 2FA

iDRAC can be configured to authenticate with a single RSA AM server at a time. The global settings on RSA AM server apply to all iDRAC local users, AD, and LDAP users.

(i) NOTE: RSA SecurID 2FA teature is available only on Datacenter license.

Following are the pre-requisites before you configure iDRAC to enable RSA SecurID:

- Configure Microsoft Active Directory server.
- If you try to enable RSA SecurID on all AD users, add the AD server to the RSA AM server as an Identity Source.
- Ensure you have a generic LDAP server.
- For all LDAP users, the Identity Source to the LDAP server must be added in RSA AM server.

To enable RSA SecurID on iDRAC, the following attributes from the RSA AM server are required:

- 1. RSA Authentication API URL The URL syntax is: https://<rsa-am-server-hostname>:<port>/mfa/v1_1, and by default the port is 5555.
- 2. **RSA Client-ID** By default, the RSA client ID is the same as the RSA AM server hostname. Find the RSA client ID at RSA AM server's authentication agent configuration page.
- 3. RSA Access Key The Access Key can be retreived on RSA AM by navigating to Setup > System Settings > RSA SecurID > Authetication APIsection, which is usually displayed as

198cv5x195fdi86u43jw0q069byt0x37umlfwxc2gnp4s0xk11ve21ffum4s8302. To configure the settings through iDRAC GUI:

- Go to **iDRAC Settings** > Users.
- From Local Users section, select an existing local user and click Edit.
- Scroll down to the bottom of the Configuration page.
- In RSA SecurID section, Click the link RSA SecurID Configuration to view or edit these settings.

You can also configure the settings as follows:

- Go to **iDRAC Settings** > Users.
- From Directory Services section, select Microsoft Active Service or Generic LDAP Directory Service, and click Edit.
- In RSA SecurID section, Click the link RSA SecurID Configuration to view or edit these settings.

4. RSA AM server certificate (chain)

You can login to iDRAC using RSA SecurID token via iDRAC GUI and SSH.

RSA SecurID Token App

You need to install RSA SecurID Token app on you system or on smart phone. When you try to log in to iDRAC, you are asked to input the passcode shown in the app.

If a wrong passcode is entered, the RSA AM server challenges the user to provide the "Next Token." This may happen even though the user may have entered the correct passcode. This entry proves that the user owns the right Token that generates the right passcode.



You get the **Next Token** from RSA SecurID Token app by clicking on **Options**. Check **Next Token**, and the next passcode is available. Time is critical in this step. Otherwise, iDRAC may fail the verification of the next token. If the iDRAC user login session times out, it requires another attempt to log in

If a wrong passcode is entered, the RSA AM server will challenge the user to provide the "Next Token." This challenge happens even though the user may have later entered the correct passcode. This entry proves that the user owns the right Token that generates the right passcodes.

To get the next token from RSA SecurID Token app, click on **Options** and check **Next Token**. A new token is generated. Time is critical in this step. Otherwise, iDRAC may fail the verification of the next token. If the iDRAC user login session times out, it requires another attempt to log in.

Viewing system health

Before you perform a task or trigger an event, you can use RACADM to check if the system is in a suitable state. To view the remote service status from RACADM, use the getremoteservicesstatus command.

Table 6. Possible values for system status

	Host System	L	ifecycle Controller (LC)		Real Time Status		Overall Status
• • • •	Powered Off In POST Out of POST Collecting System Inventory Automated Task Execution Lifecycle Controller Unified Server Configurator Server has halted at F1/F2 error prompt because of a POST error Server has halted at F1/F2/F11 prompt because there are no bootable devices available Server has entered F2 setup menu Server has entered F11 Boot Manager menu	•	Ready Not Initialized Reloading data Disabled In Recovery In Use	•	Ready Not Ready	•	Ready Not Ready
1. 2. 3. 4.	Read/Write: Read Only User Privilege: Login User License Required: iDRAC Express or iDRAC Enterprise Dependency: None						

Logging in to iDRAC using public key authentication

You can log in to the iDRAC over SSH without entering a password. You can also send a single RACADM command as a command line argument to the SSH application. The command line options behave like remote RACADM since the session ends after the command is completed.

For example:

Logging in:

ssh username@<domain>

or

ssh username@<IP_address>

Fis. <u>430</u> Mov. <u>32</u>

where IP address is the IP address of the iDRAC.

Sending RACADM commands:

ssh username@<domain> racadm getversion

```
ssh username@<domain> racadm getsel
```

Multiple iDRAC sessions

The following table provides the number of iDRAC sessions that are possible using the various interfaces.

Table 7. Multiple iDRAC sessions

Interface	Number of Sessions
iDRAC Web Interface	8
Remote RACADM	4
Firmware RACADM	SSH - 4
	Serial - 1

iDRAC allows multiple sessions for the same user. After a user has created the maximum number of allowed sessions, other users cannot log in to the iDRAC. This can cause a *Denial of Service* for a legitimate administrator user.

In case of session exhaustion, perform the following remedies:

- If webserver-based sessions are exhausted, you can still login via SSH or local RACADM.
- An administrator can then terminate existing sessions using racadm commands (racadm getssninfo; racadm closessn -i <index>).

Secure default password

All supported systems are shipped with a unique default password for iDRAC, unless you choose to set *calvin* as the password while ordering the system. The unique password helps improve the security of iDRAC and your server. To further enhance security, it is recommended that you change the default password.

The unique password for your system is available on the system information tag. To locate the tag, see the documentation for your server at https://www.dell.com/support.

(i) NOTE: For PowerEdge C6420, M640, and FC640, the default password is calvin.

NOTE: Resetting iDRAC to the factory default settings reverts the default password to the one that the server was shipped with.

If you have forgotten the password and do not have access to the system information tag, there are a few methods to reset the password locally or remotely.

Resetting default iDRAC password locally

If you have physical access to the system, you can reset the password using the following:

- iDRAC Setting utility (System Setup)
- Local RACADM
- OpenManage Mobile
- Server management USB port
- USB-NIC



Resetting default password using the iDRAC Settings utility

You can access the iDRAC settings utility using the System Setup of your server. Using the iDRAC reset to defaults all feature, you can reset the iDRAC login credentials to default.

WARNING: Resetting iDRAC to default all, resets the iDRAC to the factory defaults.

To reset iDRAC using iDRAC Settings utility:

- 1. Reboot the server and press <F2>.
- 2. In the System Setup page, click iDRAC Settings.
- 3. Click Reset iDRAC configurations to defaults all.
- 4. Click **Yes** to confirm, and then click **Back**.
- 5. Click Finish.

The server restarts after all iDRAC settings are set to default settings.

Resetting default password using local RACADM

- 1. Log in to the host OS installed on the system.
- 2. Access the local RACADM interface.
- 3. Follow the instructions in Changing the default login password using RACADM.

Resetting default password using OpenManage Mobile

You can use the OpenManage Mobile (OMM) to log in and change the default password. To log in to iDRAC using OMM, scan the QR code on the system information tag. For more information about using OMM, see the OMM documentation at *OME* - *Modular for PowerEdge MX7000 Chassis User's Guide* available at https://www.dell.com/openmanagemanuals.

them from the default values, enter the updated credentials.

Resetting default password using the server management USB port

(i) **NOTE:** These steps require that the USB management port is enabled and configured.

Using Server Configuration Profile file

Create a Server Configuration Profile (SCP) file with a new password for the default account, place it on a memory key, and use the server management USB port on the server to upload the SCP file. For more information on creating the file, see Using USB port for server management.

Accessing iDRAC using a laptop

Connect a laptop to the server management USB port and access iDRAC to change the password. For more information, see Accessing iDRAC interface over direct USB connection.

Changing default password using USB-NIC

If you have access to a keyboard, mouse, and a display device, connect to the server using the USB-NIC to access the iDRAC interface and change the default password.

- 1. Connect the devices to the system.
- 2. Use a supported browser to access the iDRAC interface using the iDRAC IP.
- **3.** Follow the instructions in Changing the default login password using web interface.



Resetting default iDRAC password remotely

If you do not have physical access to the system, you can reset the default password remotely.

Remote — Provisioned system

If you have an operating system installed on the system, use a remote desktop client to log in to the server. After you log into the server, use any of the local interfaces such as RACADM or web interface to change the password.

Remote — Non-provisioned system

If there is no operating system installed on the server and if you have a PXE setup available, use PXE and then use RACADM to reset the password.

Changing the default login password

The warning message that allows you to change the default password is displayed if:

- You log in to iDRAC with Configure User privilege.
- The default password warning feature is enabled.
- The default iDRAC user name and password are provided on the system information tag.

A warning message is also displayed when you log in to iDRAC using SSH, remote RACADM, or the Web interface. For Web interface and SSH, a single warning message is displayed for each session. For remote RACADM, the warning message is displayed for each command.

NOTE: For information on recommended characters for user names and passwords, see Recommended characters in user names and passwords.

Changing the default login password using web interface

When you log in to the iDRAC web interface, if the **Default Password Warning** page is displayed, you can change the password. To do this:

- 1. Select the Change Default Password option.
- 2. In the New Password field, enter the new password.

NOTE: For information on recommended characters for user names and passwords, see Recommended characters in user names and passwords.

- 3. In the **Confirm Password** field, enter the password again.
- 4. Click Continue.

The new password is configured and you are logged in to iDRAC.

(i) NOTE: Continue is enabled only if the passwords entered in the New Password and Confirm Password fields match.

For information about the other fields, see the *iDRAC Online Help*.

Changing the default login password using RACADM

To change the password, run the following RACADM command:

racadm set iDRAC.Users.<index>.Password <Password>

where, <index> is a value from 1 to 16 (indicates the user account) and <password> is the new user defined password. (i) NOTE: The index for the default account is 2.

For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

46 Logging in to iDRAC



NOTE: For information on recommended characters for user names and passwords, see Recommended characters in user names and passwords.

Changing the default login password using iDRAC settings utility

To change the default login password using iDRAC settings utility:

- 1. In the iDRAC Settings utility, go to User Configuration.
- The **iDRAC Settings User Configuration** page is displayed. **2.** In the **Change Password** field, enter the new password.

(i) NOTE: For information on recommended characters for user names and passwords, see Recommended characters in user names and passwords.

3. Click **Back**, click **Finish**, and then click **Yes**. The details are saved.

Enabling or disabling default password warning message

You can enable or disable the display of the default password warning message. To do this, you must have Configure Users privilege.

Password Strength Policy

Using iDRAC interface, you can check the password strength policy and check any errors if the policy is not met. The password policy cannot be applied to previously saved passwords, Server Configuration Profiles (SCP) copied from other servers, and embedded passwords in the profile.

To access Password settings, go to iDRAC Settings > Users > Password Settings.

Following fields are available in this section:

- Minimum Score Specifies the minimum password strength policy score. The values in this field are:
 - \circ 0 No protection
 - 1 Weak protection
 - 2 Medium protection
 - 3 Strong protection
- Simple Policy Specifies the required characters in a secure password. It has the following options:
 - Upper Case Letters
 - Numbers
 - Symbols
 - Minimum Length
- **Regular Expression** The Regular expression along with the Minimum score is used for password enforcement. The values are 1-4.

IP Blocking

You can use IP blocking to dynamically determine when excessive login failures occur from an IP address and block or prevent the IP address from logging into the iDRAC9 for a preselected time span. IP blocking includes:

- The number of allowable login failures.
- The timeframe in seconds when these failures must occur.
- The amount of time, in seconds, when the IP address is prevented from establishing a session after the total allowable number of failures is exceeded.



As consecutive login failures accumulate from a specific IP address, they are tracked by an internal counter. When the user logs in successfully, the failure history is cleared and the internal counter is reset.

NOTE: When consecutive login attempts are refused from the client IP address, some SSH clients may display the following message:

ssh exchange identification: Connection closed by remote host

(i) NOTE: IP blocking feature supports upto 5 IP ranges. You can see / set these only via RACADM.

Table 8. Login Retry Restriction Properties

Property	Definition
iDRAC.IPBlocking.BlockEnable	Enables the IP blocking feature. When consecutive failures
	iDRAC.IPBlocking.FailCount
	from a single IP address are encountered within a specific amount of time
	iDRAC.IPBlocking.FailWindow
	all further attempts to establish a session from that address are rejected for a certain timespan
	iDRAC.IPBlocking.PenaltyTime
iDRAC.IPBlocking.FailCount	Sets the number of login failures from an IP address before the login attempts are rejected.
iDRAC.IPBlocking.FailWindow	The time, in seconds during which the failed attempts are counted. When the failures occur beyond this time period, the counter gets reset.
iDRAC.IPBlocking.PenaltyTime	Defines the timespan in seconds when all login attempts from an IP address with excessive failures are rejected.

Enabling or disabling OS to iDRAC Pass-through using web interface

To enable OS to iDRAC Pass-through using Web interface:

- 1. Go to iDRAC Settings > Connectivity > Network > OS to iDRAC Pass-through. The OS to iDRAC Pass-through page is displayed.
- 2. Change the State to **Enabled**.
- **3.** Select any of the following options for Pass-through Mode:
 - LOM The OS to iDRAC pass-through link between the iDRAC and the host operating system is established through the LOM or NDC.
 - **USB NIC** The OS to iDRAC pass-through link between the iDRAC and the host operating system is established through the internal USB bus.
 - (i) **NOTE:** If you set the pass-through mode to LOM, ensure that:
 - OS and iDRAC are on the same subnet
 - NIC selection in Network Settings is set to LOM
- 4. If the server is connected in shared LOM mode, then the OS IP Address field is disabled.



NOTE: If the VLAN is enabled on the iDRAC, the LOM-Passthrough will only function in shared LOM mode with VLAN tagging configured on the host.

() NOTE:

- When Pass-through mode is set to LOM, it is not possible to launch iDRAC from host OS after cold boot.
- We have purposefully removed the LOM Pass-through using Dedicated mode feature.
- 5. If you select **USB NIC** as the pass-through configuration, enter the IP address of the USB NIC.

The default value is 169.254.1.1. It is recommended to use the default IP address. However, if this IP address conflicts with an IP address of other interfaces of the host system or the local network, you must change it.

Do not enter 169.254.0.3 and 169.254.0.4 IPs. These IPs are reserved for the USB NIC port on the front panel when an A/A cable is used.

- **NOTE:** If IPv6 is preferred, the default address is fde1:53ba:e9a0:de11::1. If needed, this address can be modified in the idrac.OS-BMC.UsbNicULA setting. If IPv6 is not wanted on the USB-NIC, it can be disabled by changing the address to "::"
- 6. Click Apply.
- 7. Click **Test Network Configuration** to check if the IP is accessible and the link is established between the iDRAC and the host operating system.

Enabling or disabling alerts using RACADM

Use the following command:

```
racadm set iDRAC.IPMILan.AlertEnable <n>
```

n=0 — Disabled

n=1 — Enabled

Logging in to iDRAC 49



Setting up managed system

If you need to run local RACADM or enable Last Crash Screen capture, install the following from the *Dell Systems Management Tools and Documentation* DVD:

- Local RACADM
- Server Administrator

For more information about Server Administrator, see *OpenManage Server Administrator User's Guide* available at https://www.dell.com/openmanagemanuals.

Topics:

- Setting up iDRAC IP address
- Modifying local administrator account settings
- Setting up managed system location
- Optimizing system performance and power consumption
- Setting up management station
- Configuring supported web browsers
- Updating device firmware
- Viewing and managing staged updates
- Rolling back device firmware
- Monitoring iDRAC using other Systems Management tools
- Support Server Configuration Profile Import and Export
- Secure Boot Configuration from BIOS Settings or F2
- BIOS recovery

Setting up iDRAC IP address

You must configure the initial network settings based on your network infrastructure to enable the communication to and from iDRAC. You can set up the IP address using one of the following interfaces:

- iDRAC Settings utility
- Lifecycle Controller (see Lifecycle Controller User's Guide)
- Chassis or Server LCD panel (see the system's Installation and Service Manual)
 - **NOTE:** On blade servers, you can configure the network settings using the chassis LCD panel only during the initial CMC configuration. You cannot reconfigure iDRAC using the Chassis LCD panel after the chassis is deployed.
 - CMC Web interface (not applicable for MX platforms) (see Chassis Management Controller User's Guide)

In case of rack and tower servers, you can set up the IP address or use the default iDRAC IP address 192.168.0.120 to configure initial network settings, including setting up DHCP or the static IP for iDRAC.

In case of blade servers, the iDRAC network interface is disabled by default.

After you configure iDRAC IP address:

- Ensure that you change the default user name and password.
- Access iDRAC through any of the following interfaces:
 - iDRAC Web interface using a supported browser (Internet Explorer, Firefox, Chrome, or Safari)
 - Secure Shell (SSH) Requires a client such as PuTTY on Windows. SSH is available by default in most of the Linux systems and hence does not require a client.
 - IPMITool (uses IPMI command) or shell prompt (requires Dell customized installer in Windows or Linux, available from *Systems Management Documentation and Tools* DVD or https://www.dell.com/support)



Setting up iDRAC IP using iDRAC settings utility

To set up the iDRAC IP address:

- 1. Turn on the managed system.
- 2. Press <F2> during Power-on Self-test (POST).
- 3. In the System Setup Main Menu page, click iDRAC Settings. The iDRAC Settings page is displayed.
- 4. Click Network. The Network page is displayed.
- **5.** Specify the following settings:
 - Network Settings
 - Common Settings
 - IPv4 Settings
 - IPv6 Settings
 - IPMI Settings
 - VLAN Settings
- 6. Click **Back**, click **Finish**, and then click **Yes**. The network information is saved and the system reboots.

Configuring the network settings

To configure the network settings:

(i) NOTE: For information about the options, see the *iDRAC* Settings Utility Online Help.

- 1. Under Enable NIC, select Enabled.
- 2. From the NIC Selection drop-down menu, select one of the following ports based on the network requirement:

(i) NOTE: This option is not available in MX platforms.

• **Dedicated** — Enables the remote access device to use the dedicated network interface available on the Remote Access Controller (RAC). This interface is not shared with the host operating system and routes the management traffic to a separate physical network, enabling it to be separated from the application traffic.

This option implies that iDRAC's dedicated network port routes its traffic separately from the server's LOM or NIC ports. The Dedicated option allows iDRAC to be assigned an IP address from the same subnet or different subnet in comparison to the IP addresses assigned to the Host LOM or NICs to manage the network traffic.

(i) **NOTE:** In blade servers, the Dedicated option is displayed as **Chassis (Dedicated)**.

- LOM1
- LOM2
- LOM3
- LOM4
- () NOTE: In the case of rack and tower servers, two LOM options (LOM1 and LOM2) or all four LOM options are available depending on the server model. In blade servers with two NDC ports, two LOM options (LOM1 and LOM2) are available and on servers with four NDC ports, all four LOM options are available.
- **NOTE:** Shared LOM is not supported on the *Intel 2P X520–k bNDC 10 G* if they are used in a full-height server with two NDCs because they do not support hardware arbitration.
- **3.** From the **NIC Selection** drop-down menu, select the port from which you want to access the system remotely, following are the options:
 - (i) NOTE: This feature is not available in MX platforms.
 - **NOTE:** You can select either the dedicated network interface card or from a list of LOMs available in the Quad port or Dual port mezzanine cards.



• **Chassis (Dedicated)**: Enables the remote access device to use the dedicated network interface available on the Remote Access Controller (RAC). This interface is not shared with the host operating system and routes the management traffic to a separate physical network, enabling it to be separated from the application traffic.

This option implies that iDRAC's dedicated network port routes its traffic separately from the server's LOM or NIC ports. The Dedicated option allows iDRAC to be assigned an IP address from the same subnet or different subnet in comparison to the IP addresses assigned to the Host LOM or NICs to manage the network traffic.

• For Quad port cards—LOM1-LOM16

• For Dual port cards—LOM1, LOM2, LOM5, LOM6, LOM9, LOM10, LOM13, LOM14.

4. From the **Failover Network** drop-down menu, select one of the remaining LOMs. If a network fails, the traffic is routed through the failover network.

For example, to route the iDRAC network traffic through LOM2 when LOM1 is down, select **LOM1** for **NIC Selection** and **LOM2** for **Failover Network**.

i NOTE: This option is disabled if NIC Selection is set to Dedicated.

NOTE: When using the **Failover network** settings, it is recommended that all the LOM ports to be connected to the same network.

For more details, refer to the section Modifying network settings using web interface

- Under Auto Negotiation, select On if iDRAC must automatically set the duplex mode and network speed.
 This option is available only for dedicated mode. If enabled, iDRAC sets the network speed to 10, 100, or 1000 Mbps based on the network speed.
- 6. Under Network Speed, select either 10 Mbps or 100 Mbps.
 - **NOTE:** You cannot manually set the Network Speed to 1000 Mbps. This option is available only if **Auto Negotiation** option is enabled.
- 7. Under Duplex Mode, select Half Duplex or Full Duplex option.

(i) NOTE: This option is disabled if Auto Negotiation is set to Enabled.

- () NOTE: If network teaming is configured for the host OS using the same network adapter as NIC Selection, then the Failover Network should also be configured. NIC Selection and Failover Network should use the ports that are configured as a part of the network team. If more than two ports are used as part of the network team, then the Failover Network selection should be "All".
- 8. Under NIC MTU, enter the Maximum Transmission Unit size on the NIC.
 - () NOTE: The default and maximum limit for MTU on NIC is 1500, and the minimum value is 576. An MTU value of 1280 or greater is required if IPv6 is enabled.

Common settings

If network infrastructure has DNS server, register iDRAC on the DNS. These are the initial settings requirements for advanced features such as Directory services—Active Directory or LDAP, Single Sign On, and smart card.

To register iDRAC:

- 1. Enable Register DRAC on DNS.
- 2. Enter the DNS DRAC Name.
- 3. Select Auto Config Domain Name to automatically acquire domain name from DHCP. Else, provide the DNS Domain Name.

For **DNS iDRAC Name** field, the default name format is *idrac-Service_Tag*, where Service_Tag is the service tag of the server. The maximum length is 63 characters and the following characters are supported:

- A-Z
- a-z
- 0-9
- Hyphen (-)



Configuring the IPv4 settings

To configure the IPv4 settings:

1. Select Enabled option under Enable IPv4.

(i) **NOTE:** In the 14th generation of the PowerEdge servers, DHCP is enabled by default.

- 2. Select **Enabled** option under **Enable DHCP**, so that DHCP can automatically assign the IP address, gateway, and subnet mask to iDRAC. Else, select **Disabled** and enter the values for:
 - Static IP Address
 - Static Gateway
 - Static Subnet Mask
- Optionally, enable Use DHCP to obtain DNS server address, so that the DHCP server can assign the Static Preferred DNS Server and Static Alternate DNS Server. Else, enter the IP addresses for Static Preferred DNS Server and Static Alternate DNS Server.

Configuring the IPv6 settings

Based on the infrastructure setup, you can use IPv6 address protocol.

To configure the IPv6 settings:

- **NOTE:** If IPv6 is set to static, ensure that you configure the IPv6 gateway manually, which is not needed in case of dynamic IPV6. Failing to configure manually in case of static IPV6 results in loss of communication.
- 1. Select Enabled option under Enable IPv6.
- 2. For the DHCPv6 server to automatically assign the IP address, gateway, and subnet mask to iDRAC, select **Enabled** option under **Enable Auto-configuration**.

(i) NOTE: You can configure both static IP and DHCP IP at the same time.

- 3. In the Static IP Address 1 box, enter the static IPv6 address.
- 4. In the Static Prefix Length box, enter a value between 0 and 128.
- 5. In the Static Gateway box, enter the gateway address.
 - (i) NOTE: If you configure static IP, the current IP address 1 displays static IP and the IP address 2 displays dynamic IP. If you clear the static IP settings, the current IP address 1 displays dynamic IP.
- 6. If you are using DHCP, enable **DHCPv6 to obtain DNS Server addresses** to obtain Primary and Secondary DNS server addresses from DHCPv6 server. You can configure the following if required:
 - In the Static Preferred DNS Server box, enter the static DNS server IPv6 address.
 - In the Static Alternate DNS Server box, enter the static alternate DNS server.

Configuring the IPMI settings

To enable the IPMI Settings:

- 1. Under Enable IPMI Over LAN, select Enabled.
- 2. Under Channel Privilege Limit, select Administrator, Operator, or User.
- **3.** In the **Encryption Key** box, enter the encryption key in the format 0 to 40 hexadecimal characters (without any blanks characters.) The default value is all zeros.

VLAN settings

You can configure iDRAC into the VLAN infrastructure. To configure VLAN settings, perform the following steps:

- (i) **NOTE:** On blade servers that are set as **Chassis (Dedicated)**, the VLAN settings are read-only and can be changed only using CMC. If the server is set in shared mode, you can configure VLAN settings in shared mode in iDRAC.
- 1. Under Enable VLAN ID, select Enabled.



- 2. In the VLAN ID box, enter a valid number from 1 to 4094.
- 3. In the **Priority** box, enter a number from 0 to 7 to set the priority of the VLAN ID.

(i) NOTE: After enabling VLAN, the iDRAC IP is not accessible for some time.

Setting up iDRAC IP using the CMC web interface

To set up the iDRAC IP address using the Chassis Management Controller (CMC) Web interface:

NOTE: You must have Chassis Configuration Administrator privilege to set up iDRAC network settings from CMC. The CMC option is applicable only for blade servers.

- 1. Log in to the CMC Web interface.
- 2. Go to iDRAC Settings > Settings > CMC. The Deploy iDRAC page is displayed.
- 3. Under **iDRAC Network Settings**, select **Enable LAN** and other network parameters as per requirements. For more information, see *CMC online help*.
- For additional network settings specific to each blade server, go to Server Overview > <server name>. The Server Status page is displayed.
- 5. Click Launch iDRAC and go to iDRAC Settings > Connectivity > Network.
- 6. In the **Network** page, specify the following settings:
 - Network Settings
 - Common Settings
 - IPV4 Settings
 - IPV6 Settings
 - IPMI Settings
 - VLAN Settings
 - Advanced Network Settings

(i) NOTE: For more information, see *iDRAC Online Help*.

7. To save the network information, click Apply.

For more information, see the Chassis Management Controller User's Guide available at https://www.dell.com/cmcmanuals.

Auto-discovery

The Auto-discovery feature allows newly installed servers to automatically discover the remote management console that hosts the provisioning server. The provisioning server provides custom administrative user credentials to iDRAC so that the unprovisioned server can be discovered and managed from the management console. For more information about provisioning server, see the *Lifecycle Controller Remote Services Quick Start Guide* available at https://www.dell.com/idracmanuals.

Provisioning server works with a static IP address. Auto-discovery feature on the iDRAC is used to find the provisioning server using DHCP/Unicast DNS/mDNS.

- When iDRAC has the console address, it sends its own service tag, IP address, Redfish port number, Web certificate etc.
- This information is periodically published to consoles.

DHCP, DNS server, or the default DNS host name discovers the provisioning server. If DNS is specified, the provisioning server IP is retrieved from DNS and the DHCP settings are not required. If the provisioning server is specified, discovery is skipped so neither DHCP nor DNS is required.

Auto-discovery can be enabled using the following ways:

1. Using iDRAC GUI: iDRAC Settings > Connectivity > iDRAC Auto Discover



2. Using RACADM:

🎒 jan@cobb:-				
[jontcobb -15 ssh root@ root@l0.86.059's pass /admin-> racadm get id Hkeysidra: Embeddel 14% Enable10ChamgeAnnouncer Enable10ChamgeAnnouncef Enable10ChamgeAnnouncef (Jadmin-> /admin-> fanble10ChamgeAnnounce (Jadmin-> /admin-> /admin-> fanble10ChamgeAnnounce Vaage Required License Bependency	10.30.5.30 ord: Tex.antoiscoury tradition-pailed tradition-pai			
EnableIPChangeAnnounceF Usage Required License Dependency	rodeRP Exable iDBack to obtain list of consoles through DHCP. 0 - Disabled i - Enabled Auto Discovery None			
EnableIPChangeAnnounceF Usage Required License Dependency	rommENE Exable 100AC to obtain list of consoles through mONS 0 - Disabled Enabled Noto Discovery None			
EnableIPChangeAnnounceF Usage Required License Dependency	reducicatINS Enable 1084(to obtain list of consoles through unicast DNS. 0 - 0 totabled :- Enabled Auto Discovery None			
UnsolicitedIPChangeAnno Usage Required License Dependency	unceNtate - Rate of periodic refresh of ID address to concoles - 0. Disabled 1. 1 hour; 2- 6 hours; 3- 12 hours; 4- 1 day; 5- 3 days; 6- 1 week; 7- 2 weeks; 8- 4 weeks; 6 - Auto Discovery - None	1- 6 W	reeks	
/admin1-> 📕				



To enable provisioning server using iDRAC Settings utility:

- **1.** Turn on the managed system.
- During POST, press F2, and go to iDRAC Settings > Remote Enablement. The iDRAC Settings Remote Enablement page is displayed.
- 3. Enable auto-discovery, enter the provisioning server IP address, and click **Back**.

() **NOTE:** Specifying the provisioning server IP is optional. If it is not set, it is discovered using DHCP or DNS settings (step 7).

4. Click Network.

The iDRAC Settings Network page is displayed.

- 5. Enable NIC.
- 6. Enable IPv4.

(i) NOTE: IPv6 is not supported for auto-discovery.

7. Enable DHCP and get the domain name, DNS server address, and DNS domain name from DHCP.

(i) NOTE: Step 7 is optional if the provisioning server IP address (step 3) is provided.

Configuring servers and server components using Auto Config

The Auto Config feature configures and provisions all the components in a server in a single operation. These components include BIOS, iDRAC, and PERC. Auto Config automatically imports a Server Configuration Profile (SCP) XML or JSON file containing all configurable parameters. The DHCP server that assigns the IP address also provides the details for accessing the SCP file.

SCP files are created by configuring a gold configuration server. This configuration is then exported to a shared NFS, CIFS, HTTP or HTTPS network location that is accessible by the DHCP server and the iDRAC of the server being configured. The SCP file name can be based on the Service Tag or model number of the target server or can be given as a generic name. The DHCP server uses a DHCP server option to specify the SCP file name (optionally), SCP file location, and the user credentials to access the file location.

When the iDRAC obtains an IP address from the DHCP server that is configured for Auto Config, iDRAC uses the SCP to configure the server's devices. Auto Config is invoked only after the iDRAC gets its IP address from the DHCP server. If it does not get a response or an IP address from the DHCP server, then Auto Config is not invoked.

HTTP and HTTPS file sharing options are supported for iDRAC firmware 3.00.00.00 or later. Details of the HTTP or HTTPS address need to be provided. In case the proxy is enabled on the server, the user needs to provide further proxy settings to allow HTTP or HTTPS to transfer information. The –s option flag is updated as:

Table 9. Different Share Types and pass in values

-s (ShareType)	pass in
NFS	0 or nfs
CIFS	2 or cifs
HTTP	5 or http
HTTPS	6 or https

(i) NOTE: HTTPS certificates are not supported with Auto Config. Auto Config ignores certificate warnings.

Following list describes the required and optional parameters to pass in for the string value:

-f (Filename): name of exported Server Configuration Profile file. This is required for iDRAC firmware versions prior to 2.20.20.20.

-n (Sharename): name of network share. This is required for NFS or CIFS.

-s (ShareType): pass in either 0 for NFS, 2 for CIFS, 5 for HTTP and 6 for HTTPS. This is a mandatory field for iDRAC firmware versions 3.00.00.00.

-i (IPAddress): IP address of the network share. This is a mandatory field.

56 Setting up managed system



- -u (Username): username that has access to network share. This is a mandatory field for CIFS.
- -p (Password): user password that has access to network share. This is a mandatory field for CIFS.
- -d (ShutdownType): either 0 for graceful or 1 for forced (default setting: 0). This is an optional field.
- -t (Timetowait): time to wait for the host to shutdown (default setting: 300). This is an optional field.
- -e (EndHostPowerState): either 0 for OFF or 1 for ON (default setting 1). This is an optional field.

The additional option flags are supported in iDRAC firmware 3.00.00.00 or later to enable the configuration of HTTP proxy parameters and set the retry timeout for accessing the Profile file:

- -pd (ProxyDefault): Use default proxy setting. This is an optional field.
- -pt (ProxyType): The user can pass in http or socks (default setting http). This is an optional field.
- -ph (ProxyHost): IP address of the proxy host. This is an optional field.
- -pu (ProxyUserName): username that has access to the proxy server. This is required for proxy support.
- -pp (ProxyPassword): user password that has access to the proxy server. This is required for proxy support.
- -po (ProxyPort): port for the proxy server (default setting is 80). This is an optional field.
- -to (Timeout): specifies the retry timeout in minutes for obtaining config file (default is 60 minutes).

For iDRAC firmware 3.00.00.00 or later, JSON format Profile files are supported. The following file names will be used if the Filename parameter is not present:

- <service tag>-config.xml, Example: CDVH7R1-config.xml
- <model number>-config.xml, Example: R640-config.xml
- config.xml
- <service tag>-config.json, Example:CDVH7R1-config.json
- <model number>-config.json, Example: R630-config.json
- config.json

NOTE: More information about HTTP can be found in the 14G Support for HTTP and HTTPS across IDRAC9 with Lifecycle Controller Interfaces white paper at https://www.dell.com/support.

(i) NOTE:

- Auto Config can only be enabled when DHCPv4 and the Enable IPV4 options are enabled.
- Auto Config and Auto Discovery features are mutually exclusive. Disable Auto Discovery for Auto Config to work.
- The Auto Config is disabled after a server has carried out an Auto Config operation.

If all the Dell PowerEdge servers in the DHCP server pool are of the same model type and number, then a single SCP file (config.xml) is required. The config.xml file name is used as the default SCP file name. In addition to .xml file, .json files can also be used with 14G systems. The file can be config.json.

The user can configure individual servers requiring different configuration files mapped using individual server Service Tags or server models. In an environment that has different servers with specific requirements, different SCP file names can be used to distinguish each server or server type. For example, if there are two server models to configure — PowerEdge R740s and PowerEdge R540s, use two SCP files, R740-config.xml and R540-config.xml.

NOTE: iDRAC server configuration agent automatically generates the configuration filename using the server Service Tag, model number, or the default filename — config.xml.

NOTE: If none of these files are on the network share, then the server configuration profile import job is marked as failed for file not found.

Auto Config sequence

- 1. Create or modify the SCP file that configures the attributes of Dell servers.
- Place the SCP file in a share location that is accessible by the DHCP server and all the Dell servers that are assigned IP address from the DHCP server.
- **3.** Specify the SCP file location in vendor-option 43 field of DHCP server.
- 4. The iDRAC while acquiring IP address advertises vendor class identifier. (Option 60)



- 5. The DHCP server matches the vendor class to the vendor option in the dhcpd.conf file and sends the SCP file location and, if specified the SCP file name to the iDRAC.
- 6. The iDRAC processes the SCP file and configures all the attributes listed in the file.

DHCP options

DHCPv4 allows many globally defined parameters to be passed to the DHCP clients. Each parameter is known as a DHCP option. Each option is identified with an option tag, which is a 1-byte value. Option tags 0 and 255 are reserved for padding and end of options, respectively. All other values are available for defining options.

The DHCP Option 43 is used to send information from the DHCP server to the DHCP client. The option is defined as a text string. This text string is set to contain the values of the SCP filename, share location and the credentials to access the location. For example,

```
option myname code 43 = text;
subnet 192.168.0.0 netmask 255.255.255.0 {
    # default gateway
        option routers 192.168.0.1;
    option subnet-mask 255.255.255.0;
    option nis-domain "domain.org";
    option domain-name "domain.org";
    option domain-name-servers 192.168.1.1;
    option domain-name-servers 192.168.1.1;
    option time-offset -18000; #Eastern Standard Time
    option vendor-class-identifier "iDRAC";
    set vendor-string = option vendor-class-identifier;
    option myname "-f system_config.xml -i 192.168.0.130 -u user -p password -n cifs -s
2 -d 0 -t 500";
```

where, -i is the location of the Remote File Share and -f is the file name in the string along with the credentials to the Remote File Share.

The DHCP Option 60 identifies and associates a DHCP client with a particular vendor. Any DHCP server configured to take action based on a client's vendor ID should have Option 60 and Option 43 configured. With Dell PowerEdge servers, the iDRAC identifies itself with vendor ID: *iDRAC*. Therefore, you must add a new 'Vendor Class' and create a 'scope option' under it for 'code 60,' and then enable the new scope option for the DHCP server.

Configuring option 43 on Windows

To configure option 43 on Windows:

- 1. On the DHCP server, go to **Start** > **Administration Tools** > **DHCP** to open the DHCP server administration tool.
- 2. Find the server and expand all items under it.
- **3.** Right-click on **Scope Options** and select **Configure Options**. The **Scope Options** dialog box is displayed.
- 4. Scroll down and select 043 Vendor Specific Info.
- 5. In the **Data Entry** field, click anywhere in the area under **ASCII** and enter the IP address of the server that has the share location, which contains the SCP file.

The value appears as you type it under the ASCII, but it also appears in binary to the left.

6. Click **OK** to save the configuration.

Configuring option 60 on Windows

To configure option 60 on Windows:

- 1. On the DHCP server, go to Start > Administration Tools > DHCP to open the DHCP server administration tool.
- 2. Find the server and expand the items under it.
- 3. Right-click on IPv4 and choose Define Vendor Classes.
- 4. Click Add.
 - A dialog box with the following fields is displayed:
 - Display name:
 - Description:
 - ID: Binary: ASCII:



- 5. In the **Display name:** field, type iDRAC.
- 6. In the **Description:** field, type Vendor Class.
- 7. Click in the **ASCII:** section and type iDRAC.
- 8. Click OK and then Close.
- 9. On the DHCP window, right-click IPv4 and select Set Predefined Options.
- 10. From the Option class drop-down menu, select iDRAC (created in step 4) and click Add.
- 11. In the **Option Type** dialog box, enter the following information:
 - Name iDRAC
 - Data Type String
 - Code 060
 - Description Dell vendor class identifier
- 12. Click OK to return to the DHCP window.
- 13. Expand all items under the server name, right-click Scope Options and select Configure Options.
- 14. Click the Advanced tab.
- 15. From the Vendor class drop-down menu, select **iDRAC**. The 060 iDRAC is displayed in the Available Options column.
- 16. Select 060 iDRAC option.
- 17. Enter the string value that must be sent to the iDRAC (along with a standard DHCP provided IP address). The string value helps in importing the correct SCP file.

For the option's **DATA entry, String Value** setting, use a text parameter that has the following letter options and values:

- Filename (-f) Indicates the name of the exported Server Configuration Profile(SCP) file.
- Sharename (-n) Indicates the name of the network share.
- ShareType (-s) —

Alongside supporting NFS and CIFS-based file sharing, iDRAC firmware 3.00.00.00 or later also supports accessing profile files by using HTTP and HTTPS. The -s option flag is updated as follows:

-s (ShareType): type nfs or 0 for NFS; cifs or 2 for CIFS; http or 5 for HTTP; or https or 6 for HTTPS (mandatory).

- IPAddress (-i) Indicates the IP address of the file share.
 NOTE: Sharename (-n), ShareType (-s), and IPAddress (-i) are required attributes that must be passed. -n is not required for HTTP or HTTPs.
- Username (-u) Indicates the user name required to access the network share. This information is required only for CIFS.
- Password (-p) Indicates the password required to access the network share. This information is required only for CIFS.
- ShutdownType (-d) Indicates the mode of shutdown. 0 indicates Graceful shutdown and 1 indicates Forced shutdown.

i NOTE: The default setting is 0.

- Timetowait (-t) Indicates the time the host system waits before shutting down. The default setting is 300.
- EndHostPowerState (-e) Indicates the power state of the host. 0 indicates OFF and 1 indicates ON. The default setting is 1.

(i) NOTE: ShutdownType (-d), Timetowait (-t), and EndHostPowerState (-e) are optional attributes.

NFS: -f system_config.xml -i 192.168.1.101 -n /nfs_share -s 0 -d 1 **CIFS**: -f system_config.xml -i 192.168.1.101 -n cifs_share -s 2 -u *<USERNAME>* -p *<PASSWORD>* -d 1 -t 400 **HTTP**: -f system_config.json -i 192.168.1.101 -s 5 **HTTP**: -f http_share/system_config.xml -i 192.168.1.101 -s http **HTTP**: -f system_config.xml -i 192.168.1.101 -s http -n http_share **HTTPS**: -f system_config.json -i 192.168.1.101 -s http

Configuring option 43 and option 60 on Linux

Update the /etc/dhcpd.conf file. The steps to configure the options are similar to the steps for Windows:

1. Set aside a block or pool of addresses that this DHCP server can allocate.



2. Set the option 43 and use the name vendor class identifier for option 60.

```
option myname code 43 = text;
subnet 192.168.0.0 netmask 255.255.0.0 {
#default gateway
    option routers
                                  192.168.0.1;
                                   255.255.255.0;
    option subnet-mask
                                 "domain.org";
    option nis-domain
    option domain-name
                                 "domain.org";
    option domain-name-servers
                                       192.168.1.1;
    option time-offset
                                   -18000;
                                               # Eastern Standard Time
    option vendor-class-identifier "iDRAC";
  set vendor-string = option vendor-class-identifier;
option myname "-f system_config.xml -i 192.168.0.130 -u user -p password -n cifs -s 2 -d 0 -t 500";
    range dynamic-bootp 192.168.0.128 192.168.0.254;
    default-lease-time 21600;
    max-lease-time 43200;
         }
}
```

The following are the required and optional parameters that must be passed in the vendor class identifier string:

• Filename (-f) — Indicates the name of the exported Server Configuration Profile file.

(i) NOTE: For more information on file naming rules, see Configuring servers and server components using Auto Config.

- Sharename (-n) Indicates the name of the network share.
- ShareType (-s) Indicates the share type. 0 indicates NFS, 2 indicates CIFS, 5 indicates HTTP, and 6 indicates HTTPS.
 NOTE: Example for Linux NFS, CIFS, HTTP, HTTPS share:
 - NFS:-f system config.xml -i 192.168.0.130 -n /nfs -s 0 -d 0 -t 500

Ensure that you use NFS2 or NFS3 for NFS network share.

- CIFS: -f system_config.xml -i 192.168.0.130 -n sambashare/config_files -s 2 -u user -p password -d 1 -t 400
- HTTP:-f system_config.xml -i 192.168.1.101 -s http -n http_share
- HTTPS: -f system config.json -i 192.168.1.101 -s https
- IPAddress (-i) Indicates the IP address of the file share.

NOTE: Sharename (-n), ShareType (-s), and IPAddress (-i) are required attributes that must be passed. -n is not required for HTTP or HTTPS.

- Username (-u) Indicates the user name required to access the network share. This information is required only for CIFS.
- Password (-p) Indicates the password required to access the network share. This information is required only for CIFS.
- ShutdownType (-d) Indicates the mode of shutdown. 0 indicates Graceful shutdown and 1 indicates Forced shutdown.
 (i) NOTE: The default setting is 0.
- Timetowait (-t) Indicates the time the host system waits before shutting down. The default setting is 300.
- EndHostPowerState (-e) Indicates the power state of the host. 0 indicates OFF and 1 indicates ON. The default setting is 1.

(i) NOTE: ShutdownType (-d), Timetowait (-t), and EndHostPowerState (-e) are optional attributes.

The following is an example of a static DHCP reservation from a dhcpd.conf file:

```
host my_host {
host my_host {
hardware ethernet b8:2a:72:fb:e6:56;
fixed-address 192.168.0.211;
option host-name "my_host";
option myname " -f r630_raid.xml -i 192.168.0.1 -n /nfs -s 0 -d 0 -t 300";
}
```

(i) NOTE: After editing the dhcpd.conf file, make sure to restart the dhcpd service to apply the changes.



Prerequisites before enabling Auto Config

Before enabling the Auto config feature, make sure that following are already set:

- Supported network share (NFS, CIFS, HTTP and HTTPS) is available on the same subnet as the iDRAC and DHCP server. Test the network share to ensure that it can be accessed and that the firewall and user permissions are set correctly.
- Server configuration profile is exported to the network share. Also, make sure that the necessary changes in the SCP file are complete so that proper settings can be applied when the Auto Config process is initiated.
- DHCP server is set and the DHCP configuration is updated as required for iDRAC to call the server and initiate the Auto Config feature.

Enabling Auto Config using iDRAC web interface

Make sure that DHCPv4 and the Enable IPv4 options are enabled and Auto-discovery is disabled.

To enable Auto Config:

- 1. In the iDRAC web interface, go to **iDRAC Settings** > **Connectivity** > **Network** > **Auto Config**. The **Network** page is displayed.
- 2. In the Auto Config section, select one of the following options from the Enable DHCP Provisioning drop-down menu:
 - Enable Once Configures the component only once using the SCP file referenced by the DHCP server. After this, Auto Config is disabled.
 - **Enable once after reset** After the iDRAC is reset, configures the components only once using the SCP file referenced by the DHCP server. After this, Auto Config is disabled.
 - **Disable** Disables the Auto Config feature.
- **3.** Click **Apply** to apply the setting. The network page automatically refreshes.

Enabling Auto Config using RACADM

To enable Auto Config feature using RACADM, use the iDRAC.NIC.AutoConfig object.

For more information, see the iDRAC RACADM CLI Guide available at https://www.dell.com/idracmanuals.

For more information on the Auto Config feature, see the Zero-Touch, bare-metal server provisioning using the Dell EMC iDRAC with Lifecycle Controller Auto Config feature white paper available at the https://www.dell.com/support.

Using hash passwords for improved security

On PowerEdge servers with iDRAC version 3.00.00.00, you can set user passwords and BIOS passwords using a one-way hash format. The user authentication mechanism is not affected (except for SNMPv3 and IPMI) and you can provide the password in plain text format.

With the new password hash feature:

• You can generate your own SHA256 hashes to set iDRAC user passwords and BIOS passwords. This allows you to have the SHA256 values in the server configuration profile, RACADM, and WSMan. When you provide the SHA256 password values, you cannot authenticate through SNMPv3 and IPMI.

NOTE: Remote RACADM or WSMan or Redfish cannot be used for Hash password Configuration / Replacement for IDRAC. You can use SCP for Hash Password Configuration / Replacement on Remote RACADM or WSMan or Redfish.

- You can set up a template server including all the iDRAC user accounts and BIOS passwords using the current plain text
 mechanism. After the server is set up, you can export the server configuration profile with the password hash values. The
 export includes the hash values required for SNMPv3 and IPMI authentication. After importing this profile, you must use the
 the latest Dell IPMI tool, if you use an older tool, the IPMI authentication fails for the users who have the hashed password
 values set.
- The other interfaces such as IDRAC GUI will show the user accounts enabled.

You can generate the hash password with and without Salt using SHA256.

You must have Server Control privileges to include and export hash passwords.

If access to all accounts is lost, use iDRAC Settings Utility or local RACADM and perform reset iDRAC to default task.

If the password of the iDRAC user account is set with the SHA256 password hash only and not the other hashes (SHA1v3Key or MD5v3Key or IPMIKey), then authentication through SNMP v3 and IPMI is not available.

Fls. <u>448</u> Mov. 32

Hash password using RACADM

To set hash passwords, use the following objects with the set command:

- iDRAC.Users.SHA256Password
- iDRAC.Users.SHA256PasswordSalt

() NOTE: SHA256Password and SHA256PasswordSalt fields are reserved for XML import and do not set them using command line tools. Setting one of the fields can potentially lock out the current user from logging into iDRAC. When a password is imported using SHA256Password, the iDRAC will not be enforcing the password length check.

Use the following command to include the hash password in the exported server configuration profile:

racadm get -f <file name> -l <NFS / CIFS / HTTP / HTTPS share> -u <username> -p
<password> -t <filetype> --includePH

You must set the Salt attribute when the associated hash is set.

(i) NOTE: The attributes are not applicable to the INI configuration file.

Hash password in server configuration profile

The new hash passwords can be optionally exported in the server configuration profile.

When importing server configuration profile, you can uncomment the existing password attribute or the new password hash attribute(s). If both are uncommented an error is generated and the password is not set. A commented attribute is not applied during an import.

Generating hash password without SNMPv3 and IPMI authentication

Hash password can be generated without SNMPv3 and IPMI authentication with or without salt. Both require SHA256.

To generate hash password with salt:

1. For the iDRAC user accounts, you must salt the password using SHA256.

When you salt the password, a 16-bytes binary string is appended. The Salt is required to be 16 bytes long, if provided. Once appended, it becomes a 32 character string. The format is "password"+"salt", for example:

Password = SOMEPASSWORD

Salt = ALITTLEBITOFSALT-16 characters are appended

2. Open a Linux command prompt, and run the following command:

```
Generate Hash-> echo-n SOMEPASSWORDALITTLEBITOFSALT|sha256sum -><HASH>
```

Generate Hex Representation of Salt -> echo -n ALITTLEBITOFSALT | xxd -p -> <HEX-SALT>

set iDRAC.Users.4.SHA256Password <HASH>

set iDRAC.Users.4.SHA256PasswordSalt <HEX-SALT>

3. Provide hash value and salt in the imported server configuration profile, the RACADM commands, Redfish, or WSMan.



NOTE: If you wish to clear a previously salted password, then ensure that the password-salt is explicitly set to an empty string i.e.

```
set iDRAC.Users.4.SHA256Password
ca74e5fe75654735d3b8d04a7bdf5dcdd06f1c6c2a215171a24e5a9dcb28e7a2
```

set iDRAC.Users.4.SHA256PasswordSalt

 After setting the password, the normal plain text password authentication works except that SNMP v3 and IPMI authentication fails for the iDRAC user accounts that had passwords updated with hash.

Modifying local administrator account settings

After setting the iDRAC IP address, you can modify the local administrator account settings (that is, user 2) using the iDRAC Settings utility. To do this:

- In the iDRAC Settings utility, go to User Configuration. The iDRAC Settings User Configuration page is displayed.
- 2. Specify the details for User Name, LAN User Privilege, Serial Port User Privilege, and Change Password. For information about the options, see the *iDRAC Settings Utility Online Help*.
- **3.** Click **Back**, click **Finish**, and then click **Yes**. The local administrator account settings are configured.

Setting up managed system location

You can specify the location details of the managed system in the data center using the iDRAC Web interface or iDRAC Settings utility.

Setting up managed system location using web interface

To specify the system location details:

- In the iDRAC web interface, go to System > Details > System Details. The System Details page is displayed.
- **2.** Under **System Location**, enter the location details of the managed system in the data center. For information about the options, see the *iDRAC Online Help*.
- 3. Click Apply. The system location details are saved in iDRAC.

Setting up managed system location using RACADM

To specify the system location details, use the System.Location group objects.

For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Setting up managed system location using iDRAC settings utility

To specify the system location details:

- In the iDRAC Settings utility, go to System Location. The iDRAC Settings System Location page is displayed.
- 2. Enter the location details of the managed system in the data center. For information about the options, see the *iDRAC* Settings Utility Online Help.
- **3.** Click **Back**, click **Finish**, and then click **Yes**. The details are saved.



Optimizing system performance and power consumption

The power required to cool a server can contribute a significant amount to the overall system power. Thermal control is the active management of system cooling through fan speed and system power management to make sure that the system is reliable while minimizing system power consumption, airflow, and system acoustic output. You can adjust the thermal control settings and optimize against the system performance and performance-per-Watt requirements.

Using the iDRAC Web interface, RACADM, or the iDRAC Settings Utility, you can change the following thermal settings:

- Optimize for performance
- Optimize for minimum power
- Set the maximum air exhaust temperature
- Increase airflow through a fan offset, if required

• Increase airflow through increasing minimum fan speed

Following are the list of features in thermal management:

- System Airflow Consumption: Displays the real-time system airflow consumption (in CFM), allowing airflow balancing at rack and datacenter level.
- Custom Delta-T: Limit air temperature rise from inlet air to exhaust to right-size your infrastructure level cooling.
- Exhaust Temperature Control: Specify the temperature limit of the air exiting the server to match your datacenter needs.
- **Custom PCIe inlet temperature**: Choose the right input inlet temperature to match 3rd party device requirements.
- **PCIe Airflow settings**: Provides a comprehensive PCIe device cooling view of the server and allows cooling customization of 3rd party cards.

Modifying thermal settings using iDRAC web interface

To modify the thermal settings:

- 1. In the iDRAC Web interface, go to Configuration > System Settings > Hardware Settings > Cooling Configuration.
- 2. Specify the following:
 - Thermal Profile Optimization Select the thermal profile:
 - **Default Thermal Profile Settings (Minimum Power)** Implies that the thermal algorithm uses the same system profile settings that is defined under **System BIOS** > **System BIOS** > **System Profile Settings** page.

By default, this option is set to **Default Thermal Profile Settings**. You can also select a custom algorithm, which is independent of the BIOS profile. The options available are:

- Maximum Performance (Performance Optimized) :
 - Reduced probability of memory or CPU throttling.
 - Increased probability of turbo mode activation.
 - Generally, higher fan speeds at idle and stress loads.
- Minimum Power (Performance per Watt Optimized):
 - Optimized for lowest system power consumption based on optimum fan power state.
 - Generally, lower fan speeds at idle and stress loads.
- **Sound Cap** Sound Cap provides reduced acoustical output from a server at the expense of some performance. Enabling Sound Cap may include temporary deployment or evaluation of a server in an occupied space, but it should not be used during benchmarking or performance sensitive applications.

(i) **NOTE:** Selecting **Maximum Performance** or **Minimum Power**, overrides thermal settings associated to System Profile setting under **System BIOS** > **System BIOS** Settings.System Profile Settings page.

• **Maximum Exhaust Temperature Limit** — From the drop-down menu, select the maximum exhaust air temperature. The values are displayed based on the system.

The default value is Default, 70°C (158 °F).

This option allows the system fans speeds to change such that the exhaust temperature does not exceed the selected exhaust temperature limit. This cannot always be guaranteed under all system operating conditions due to dependency on system load and system cooling capability.

• Fan Speed Offset — Selecting this option allows additional cooling to the server. In case hardware is added (example, new PCle cards), it may require additional cooling. A fan speed offset causes fan speeds to increase (by the offset % value) over baseline fan speeds calculated by the Thermal Control algorithm. Possible values are:


- Low Fan Speed Drives fan speeds to a moderate fan speed.
- Medium Fan Speed Drives fan speeds close to medium.
- **High Fan Speed** Drives fan speeds close to full speed.
- Max Fan Speed Drives fan speeds to full speed.
- **Off** Fan speed offset is set to off. This is the default value. When set to off, the percentage does not display. The default fan speed is applied with no offset. Conversely, the maximum setting will result in all fans running at maximum speed.

The fan speed offset is dynamic and based on the system. The fan speed increase for each offset is displayed next to each option.

The fan speed offset increases all fan speeds by the same percentage. Fan speeds may increase beyond the offset speeds based on individual component cooling needs. The overall system power consumption is expected to increase.

Fan speed offset allows you to increase the system fan speed with four incremental steps. These steps are equally divided between the typical baseline speed and the maximum speed of the server system fans. Some hardware configurations results in higher baseline fan speeds, which results in offsets other than the maximum offset to achieve maximum speed.

The most common usage scenario is non-standard PCIe adapter cooling. However, the feature can be used to increase system cooling for other purposes.

NOTE: FAN configuration setting is available in iDRAC even when system does not have any FANs. This is because, iDRAC sends the specified configuration to Chassis manager and Chassis manager can process the data from iDRAC and send the required cooling to the system as per configuration.

Thresholds

- **Maximum PCIe Inlet Temperature Limit** Default value is 55°C. Select the lower temperature of 45°C for third party PCIe cards which require lower inlet temperature.
 - **Exhaust Temperature Limits** By modifying the values for the following you can set the exhaust temperature limits:
 - Set Maximum Exhaust Temperature Limit
 - Set Air Temperature Rise Limit
- **Minimum Fan Speed in PWM (% of Max)** Select this option to fine tune the fan speed. Using this option, you can set a higher baseline system fan speed or increase the system fan speed if other custom fan speed options are not resulting in the required higher fan speeds.
 - **Default** Sets minimum fan speed to default value as determined by the system cooling algorithm.
 - **Custom** Enter the percentage by which you want to change the fan speed. Range is between 9-100.

The allowable range for minimum fan speed PWM is dynamic based on the system configuration. The first value is the idle speed and the second value is the configuration max (Depending on the system configuration, the maximum speed may be up to 100%.).

System fans can run higher than this speed as per thermal requirements of the system but not lower than the defined minimum speed. For example, setting Minimum Fan Speed at 35% limits the fan speed to never go lower than 35% PWM.

(i) NOTE: 0% PWM does not indicate fan is off. It is the lowest fan speed that the fan can achieve.

The settings are persistent, which means that once they are set and applied, they do not automatically change to the default setting during system reboot, power cycling, iDRAC, or BIOS updates. The custom cooling options may not be supported on all servers. If the options are not supported, they are not displayed or you cannot provide a custom value.

3. Click Apply to apply the settings.

The following message is displayed:

It is recommended to reboot the system when a thermal profile change has been made. This is to ensure all power and thermal settings are activated.

4. Click Reboot Later or Reboot Now.

(i) NOTE: You must reboot the system for the settings to take effect.



Modifying thermal settings using RACADM

To modify the thermal settings, use the objects in the **system.thermalsettings** group with the **set** sub command as provided in the following table.

Table 10. Thermal Settings

Object	Description	Usage	Example
AirExhaustTemp A n t	Allows you to set the maximum air exhaust temperature limit.	Set to any of the following values (based on the system): • 0 — Indicates 40°C • 1 — Indicates 45°C • 2 — Indicates 50°C • 3 — Indicates 55°C	To check the existing setting on the system: racadm get system.thermalsetti ngs.AirExhaustTemp The output is:
		 255 — Indicates 70°C (default) 	AirExhaustTemp=70
			This output indicates that the system is set to limit the air exhaust temperature to 70°C. To set the exhaust temperature limit to 60°C:
			racadm set system.thermalsetti ngs.AirExhaustTemp 4
			The output is:
			Object value modified successfully.
			If a system does not support a particular air exhaust temperature limit, then when you run the following command:
			racadm set system.thermalsetti ngs.AirExhaustTemp 0
			The following error message is displayed:
			ERROR: RAC947: Invalid object value specified.
			Make sure to specify the value depending on the type of object.
			For more information, see RACADM help.



Table 10. Thermal Settings (continued)

Object	Description	Usage	Example
			To set the limit to the default value:
			racadm set system.thermalsetti ngs.AirExhaustTemp 255
FanSpeedHighOffsetVal	 Getting this variable reads the fan speed offset value in %PWM for High Fan Speed Offset setting. This value depends on the overtam 	Values from 0-100	racadm get system.thermalsetti ngs FanSpeedHighOffsetV al
	 Use FanSpeedOffset object to set this value using index value 1. 		A numerical value, for example 66, is returned. This value indicates that when you use the following command, it applies a fan speed offset of High (66% PWM) over the baseline fan speed
			racadm set system.thermalsetti ngs FanSpeedOffset 1
FanSpeedLowOffsetVal	 Getting this variable reads the fan speed offset value in %PWM for Low Fan Speed Offset setting. This value depends on the system 	Values from 0-100	racadm get system.thermalsetti ngs FanSpeedLowOffsetVa l
 System. Use FanSpeedOffset object to set this value using index value 0. 		This returns a value such as "23". This means that when you use the following command, it applies a fan speed offset of Low (23% PWM) over baseline fan speed	
			racadm set system.thermalsetti ngs FanSpeedOffset 0
FanSpeedMaxOffsetVal	 Getting this variable reads the fan speed offset value in %PWM for Max Fan Speed Offset setting. This value depends on the system 	Values from 0-100	racadm get system.thermalsetti ngs FanSpeedMaxOffsetVa l
	 Use FanSpeedOffset to set this value using index value 3 		This returns a value such as "100". This means that when you use the following command, it applies a fan speed offset of Max (meaning full speed, 100% PWM). Usually, this offset results in



Table 10. Thermal Settings (continued)

Object	Description	Usage	Example
			fan speed increasing to full speed.
			racadm set system.thermalsetti ngs FanSpeedOffset 3
FanSpeedMediumOffsetV	Getting this variable reads	Values from 0-100	racadm get
al	the fan speed offset value in %PWM for Medium Fan Speed Offset setting.This value depends on the		system.thermalsetti ngs FanSpeedMediumOffse tVal
	 Use FanSpeedOffset object to set this value using index value 2 		This returns a value such as "47". This means that when you use the following command, it applies a fan speed offset of Medium (47% PWM) over baseline fan speed
			racadm set system.thermalsetti ngs FanSpeedOffset 2
FanSpeedOffset	Using this object with get command displays the	Values are:	To view the existing setting:
 get command displays to existing Fan Speed Offs value. Using this object with s command allows setting the required fan speed offset value. The index value decides 	existing Fan Speed Offset value.Using this object with set command allows actives	 1 — High Fan Speed 2 — Medium Fan Speed 3 — Max Fan Speed 255 — None 	racadm get system.thermalsetti ngs.FanSpeedOffset
	 the required fan speed offset value. The index value decides 		To set the fan speed offset to High value (as defined in FanSpeedHighOffsetVal)
	the offset that is applied and the FanSpeedLowOffsetVa 1, FanSpeedMaxOffsetVa		racadm set system.thermalsetti ngs.FanSpeedOffset 1
	l, FanSpeedHighOffsetV al, and FanSpeedMediumOffse tVal objects (defined earlier) are the values at which the offsets are applied.		
MFSMaximumLimit	Read Maximum limit for MFS	Values from 1 — 100	To display the highest value that can be set using MinimumFanSpeed option:
			racadm get system.thermalsetti ngs.MFSMaximumLimit

Inserido ao protocolo **22.951.206-4** por: **Pedro Henrique Golin Linhares** em: 31/03/2025 12:49. A autenticidade deste documento pode ser validada no endereço: https://www.eprotocolo.pr.gov.br/spiweb/validarDocumento com o código: **692a5ef82beb173ed33dad4ee6505546**.



Table 10. Thermal Settings (continued)

Object	Description	Usage	Example
MFSMinimumLimit	Read Minimum limit for MFS	Values from 0 to MFSMaximumLimit Default is 255 (means None)	To display the lowest value that can be set using MinimumFanSpeed option. racadm get system.thermalsetti ngs.MFSMinimumLimit
MinimumFanSpeed	 Allows configuring the Minimum Fan speed that is required for the system to operate. It defines the baseline (floor) value for fan speed and system allows fans to go lower than this defined fan speed value. This value is %PWM value for fan speed. 	Values from MFSMinimumLimit to MFSMaximumLimit When get command reports 255, it means user configured offset is not applied.	To make sure that the system minimum speed does not decrease lower than 45% PWM (45 must be a value between MFSMinimumLimit to MFSMaximumLimit): racadm set system.thermalsetti ngs.MinimumFanSpeed 45
ThermalProfile	 Allows you to specify the Thermal Base Algorithm. Allows you to set the system profile as required for thermal behavior associated to the profile. 	 Values: 0 — Auto 1 — Maximum performance 2 — Minimum Power 	To view the existing thermal profile setting: racadm get system.thermalsetti ngs.ThermalProfile To set the thermal profile to Maximum Performance: racadm set system.thermalsetti ngs.ThermalProfile 1
ThirdPartyPCIFanRespo nse	 Thermal overrides for third-party PCI cards. Allows you to disable or enable the default system fan response for detected third-party PCI cards. You can confirm the presence of third-party PCI card by viewing the message ID PCI3018 in the Lifecycle Controller log. 	Values: • 1 — Enabled • 0 — Disabled (i) NOTE: The default value is 1.	To disable any default fan speed response set for a detected third-party PCI card: racadm set system.thermalsetti ngs.ThirdPartyPCIFa nResponse 0

Modifying thermal settings using iDRAC settings utility

To modify the thermal settings:

- In the iDRAC Settings utility, go to Thermal. The iDRAC Settings Thermal page is displayed.
- 2. Specify the following:
 - Thermal Profile
 - Maximum Exhaust Temperature Limit
 - Fan Speed Offset
 - Minimum Fan Speed



The settings are persistent, which means that once they are set and applied, they do not automatically change to the default setting during system reboot, power cycling, iDRAC, or BIOS updates. A few Dell servers may or may not support some or all of these custom user cooling options. If the options are not supported, they are not displayed or you cannot provide a custom value.

3. Click **Back**, click **Finish**, and then click **Yes**. The thermal settings are configured.

Modifying PCIe airflow settings using iDRAC web interface

Use the PCIe airflow settings when increased thermal margin is desired for custom high powered PCIe cards.

To modify the PCIe airflow settings:

- 1. In the iDRAC Web interface, go to Configuration > System Settings > Hardware Settings > Cooling Configuration. The PCIe Airflow Settings page is displayed below the fan settings section.
- 2. Specify the following:
 - LFM Mode Select the Custom mode to enable the custom LFM option.
 - **Custom LFM** Enter the LFM value.
- **3.** Click **Apply** to apply the settings.

The following message is displayed:

It is recommended to reboot the system when a thermal profile change has been made. This is to ensure all power and thermal settings are activated.

Click Reboot Later or Reboot Now.

(i) NOTE: You must reboot the system for the settings to take effect.

Setting up management station

A management station is a computer used for accessing iDRAC interfaces to remotely monitor and manage the PowerEdge server(s).

To set up the management station:

- 1. Install a supported operating system. For more information, see the release notes.
- 2. Install and configure a supported Web browser. For more information, see the release notes.
- **3.** Install the latest Java Runtime Environment (JRE) (required if Java plug-in type is used to access iDRAC using a Web browser).

(i) NOTE: You need Java 8 or later to use this feature and to launch iDRAC Virtual Console over an IPv6 network.

- 4. From the *Dell Systems Management Tools and Documentation* DVD, install Remote RACADM VMCLI from the SYSMGMT folder. Else, run **Setup** on the DVD to install Remote RACADM by default and other OpenManage software. For more information about RACADM, see *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.
- 5. Install the following based on the requirement:
 - SSH client
 - TFTP
 - Dell OpenManage Essentials

Accessing iDRAC remotely

To remotely access iDRAC Web interface from a management station, make sure that the management station is in the same network as iDRAC. For example:

• Blade servers — The management station must be on the same network as CMC and OME Modular. For more information on isolating CMC network from the managed system's network, see the *Chassis Management Controller User's Guide* available at https://www.dell.com/cmcmanuals.



• Rack and tower servers — Set the iDRAC NIC to Dedicated or LOM1 and make sure that the management station is on the same network as iDRAC.

To access the managed system's console from a management station, use Virtual Console through iDRAC Web interface.

Configuring supported web browsers

() NOTE: For information about the supported browsers and their versions, see the *Release Notes* available at https://www.dell.com/idracmanuals.

Most features of iDRAC web interface can be accessed using these browsers with default settings. For certain feature to work, you must change a few settings. These settings include disabling pop-up blockers, enabling Java, ActiveX, or HTML5 plug-in support and so on.

If you are connecting to iDRAC web interface from a management station that connects to the Internet through a proxy server, configure the web browser to access the Internet through this server.

NOTE: If you use Internet Explorer or Firefox to access the iDRAC web interface, you may need to configure certain settings as described in this section. You can use other supported browsers with their default settings.

(i) NOTE: Blank proxy settings are treated equivalent to No proxy.

Configuring Internet Explorer

This section provides details about configuring Internet Explorer (IE) to ensure you can access and use all features of the iDRAC web interface. These settings include:

- Resetting security settings
- Adding iDRAC IP to trusted sites
- Configuring IE to enable Active Directory SSO
- Disabling IE Enhanced Security Configuration

Resetting Internet Explorer security settings

Ensure that Internet Explorer (IE) settings are set to Microsoft-recommended defaults and customize the settings as described in this section.

- 1. Open IE as an administrator or using an administrator account.
- 2. Click Tools Internet Options Security Local Network or Local intranet.
- 3. Click Custom Level, select Medium-Low, and click Reset. Click OK to confirm.

Adding iDRAC IP to the trusted-sites list

When you access iDRAC Web interface, you are prompted to add iDRAC IP address to the list of trusted domains if the IP address is missing from the list. When completed, click **Refresh** or relaunch the web browser to establish a connection to iDRAC web interface. If you are not prompted to add the IP, it is recommended that you add the IP manually to the trusted-sites list.

() **NOTE:** When connecting to the iDRAC web interface with a certificate the browser does not trust, the browser's certificate error warning may display a second time after you acknowledge the first warning.

To add iDRAC IP address to the trusted-sites list:

- 1. Click Tools > Internet Options > Security > Trusted sites > Sites.
- 2. Enter the iDRAC IP address to the Add this website to the zone.
- 3. Click Add, click OK, and then click Close.
- 4. Click **OK** and then refresh your browser.



Configuring Internet Explorer to enable Active Directory SSO

To configure the browser settings for Internet Explorer:

- 1. In Internet Explorer, navigate to Local Intranet and click Sites.
- **2.** Select the following options only:
- Include all local (intranet) sites not listed on other zones.
 - Include all sites that bypass the proxy server.
- 3. Click Advanced.
- Add all relative domain names that will be used for iDRAC instances that is part of the SSO configuration (for example, myhost.example.com.)
- 5. Click Close and click OK twice.

Disabling Internet Explorer Enhanced Security Configuration

To ensure that you can download log files and other local elements using the web interface, it is recommended to disable Internet Explorer Enhanced Security Configuration from Windows features. For information about disabling this feature on your version of Windows, see Microsoft's documentation.

Configuring Mozilla Firefox

This section provides details about configuring Firefox to ensure you can access and use all features of the iDRAC web interface. These settings include:

- Disabling whitelist feature
- Configuring Firefox to enable Active Directory SSO

(i) NOTE: Mozilla Firefox browser may not have scroll bar for iDRAC Online Help page.

Disabling whitelist feature in Firefox

Firefox has a "whitelist" security feature that requires user permission to install plug-ins for each distinct site that hosts a plug-in. If enabled, the whitelist feature requires you to install a Virtual Console viewer for each iDRAC you visit, even though the viewer versions are identical.

To disable the whitelist feature and avoid unnecessary plug-in installations, perform the following steps:

- 1. Open a Firefox Web browser window.
- $\ensuremath{\text{2. In the address field, enter about: config and press < Enter>}. \ensuremath{$
- 3. In the **Preference Name** column, locate and double-click **xpinstall.whitelist.required**.
- The values for **Preference Name**, **Status**, **Type**, and **Value** change to bold text. The **Status** value changes to user set and the **Value** changes to false.
- In the Preferences Name column, locate xpinstall.enabled.
 Make sure that Value is true. If not, double-click xpinstall.enabled to set Value to true.

Configuring Firefox to enable Active Directory SSO

To configure the browser settings for Firefox:

- 1. In Firefox address bar, enter about:config.
- 2. In Filter, enter network.negotiate.
- 3. Add the domain name to network.negotiate-auth.trusted-uris (using comma separated list.)
- 4. Add the domain name to network.negotiate-auth.delegation-uris (using comma separated list.)

Configuring web browsers to use virtual console

To use Virtual Console on your management station:

72 Setting up managed system



 Make sure that a supported version of the browser (Internet Explorer (Windows), or Mozilla Firefox (Windows or Linux), Google Chrome, Safari) is installed.

For more information about the supported browser versions, see the *Release Notes* available at https://www.dell.com/ idracmanuals.

- 2. To use Internet Explorer, set IE to Run As Administrator.
- 3. Configure the Web browser to use ActiveX, Java, or HTML5 plug-in.

ActiveX viewer is supported only with Internet Explorer. HTML5 or a Java viewer is supported on any browser.

INOTE: You need Java 8 or later to use this feature and to launch iDRAC Virtual Console over an IPv6 network.

- 4. Import the root certificates on the managed system to avoid the pop-ups that prompt you to verify the certificates.
- 5. Install the compat-libstdc++-33-3.2.3-61 related package.

NOTE: On Windows, the compat-libstdc++-33-3.2.3-61 related package may be included in the .NET framework package or the operating system package.

6. If you are using MAC operating system, select the **Enable access for assistive devices** option in the **Universal Access** window.

For more information, see the MAC operating system documentation.

Configuring Internet Explorer to use HTML5-based plug-in

The HTML5 virtual console and virtual media APIs are created by using HTML5 technology. The following are the advantages of HTML5 technology:

- Installation is not required on the client workstation.
- Compatibility is based on browser and is not based on the operating system or installed components.
- Compatible with most of the desktops and mobile platforms.
- Quick deployment and the client is downloaded as part of a web page.

You must configure Internet Explorer (IE) settings before you launch and run HTML5 based virtual console and virtual media applications. To configure the browser settings:

- 1. Disable pop-up blocker. To do this, click **Tools** > **Internet Options** > **Privacy** and clear the **Turn on Pop-up Blocker** check-box.
- 2. Start the HTML5 virtual console using any of the following methods:
 - In IE, click Tools > Compatibility View Settings and clear the Display intranet sites in Compatibility View checkbox.
 - In IE using an IPv6 address, modify the IPv6 address as follows:

https://[fe80::d267:e5ff:fef4:2fe9]/ to https://fe80--d267-e5ff-fef4-2fe9.ipv6literal.net/

• Direct HTML5 virtual console in IE using an IPv6 address, modify the IPv6 address as follows:

https://[fe80::d267:e5ff:fef4:2fe9]/console to https://fe80--d267-e5fffef4-2fe9.ipv6-literal.net/console

3. To display the Title Bar information in IE, go to Control Panel > Appearance and Personalization > Personalization > Window Classic

Configuring Microsoft Edge to use HTML5-based plug-in

You must configure Edge settings before you launch and run HTML5 based virtual console and virtual media applications. To configure the browser settings:

- 1. Click Settings > View Advanced Settings and disable the Block pop-ups option.
- **2.** Modify the IPv6 address as follows :

```
https://2607:f2b1:f083:147::leb.ipv6:literal.net/restgui to https://2607-f2b1-
f083-147--leb.ipv6-literal.net/restgui
```



Configuring the web browser to use Java plug-in

Install a Java Runtime Environment (JRE) if you are using Firefox or IE and want to use the Java Viewer.

NOTE: Install a 32-bit or 64-bit JRE version on a 64-bit operating system or a 32-bit JRE version on a 32-bit operating system.

To configure IE to use Java plug-in:

- Disable automatic prompting for file downloads in Internet Explorer.
- Disable Enhanced Security Mode in Internet Explorer.

Configuring IE to use ActiveX plug-in

You must configure the IE browser settings before you start and run ActiveX based Virtual Console and Virtual Media applications. The ActiveX applications are delivered as signed CAB files from the iDRAC server. If the plug-in type is set to Native-ActiveX type in Virtual console, when you try to start the Virtual Console, the CAB file is downloaded to the client system and ActiveX based Virtual Console is started. Internet Explorer requires some configurations to download, install, and run these ActiveX based applications.

On 64-bit operating systems, you can install both 32-bit or 64-bit versions of Internet Explorer. You may use either of 32-bit or 64-bit, however, you must install the corresponding plug-in. For example, if you install the plug-in in the 64-bit browser and then open the viewer in a 32-bit browser, you must install the plug-in again.

(i) NOTE: You can use ActiveX plug-in only with Internet Explorer.

NOTE: To use ActiveX plug-in on systems with Internet Explorer 9, before configuring Internet Explorer, ensure that you disable the Enhanced Security Mode in Internet Explorer or in the server manager in Windows Server operating systems.

For ActiveX applications in Windows 7, Windows 2008, and Windows 10 configure the following Internet Explorer settings to use the ActiveX plug-in:

- 1. Clear the browser's cache.
- 2. Add iDRAC IP or host name to the Local Internet site list.
- 3. Reset the custom settings to **Medium-low** or change the settings to allow installation of signed ActiveX plug-ins.
- 4. Enable the browser to download encrypted content and to enable third-party browser extensions. To do this, go to Tools > Internet Options > Advanced, clear the Do not save encrypted pages to disk option, and select the Enable third-party browser extensions option.

(i) NOTE: Restart Internet Explorer for the Enable third-party browser extension setting to take effect.

- 5. Go to Tools > Internet Options > Security and select the zone in which you want to run the application.
- 6. Click Custom level. In the Security Settings window, do the following:
 - Select Enable for Automatic prompting for ActiveX controls.
 - Select Prompt for Download signed ActiveX controls.
 - Select Enable or Prompt for Run ActiveX controls and plugins.
 - Select Enable or Prompt for Script ActiveX controls marked safe for scripting.
- 7. Click OK to close the Security Settings window.
- 8. Click OK to close the Internet Options window.
 - **NOTE:** On systems with Internet Explorer 11, ensure that you add the iDRAC IP by clicking **Tools** > **Compatibility View settings**.
 - () NOTE:
 - The varying versions of Internet Explorer share **Internet Options**. Therefore, after you add the server to the list of *trusted sites* for one browser the other browser uses the same setting.
 - Before installing the ActiveX control, Internet Explorer may display a security warning. To complete the ActiveX control installation procedure, accept the ActiveX control when Internet Explorer prompts you with a security warning.



• If you get the error **Unknown Publisher** while launching Virtual Console, it may be caused because of the change to the code-signing certificate path. To resolve this error, you must download an addition key. Use a search engine to search for **Symantec SO16958** and, from the search results, follow the instructions on the Symantec website.

Additional settings for Windows Vista or newer Microsoft operating systems

The Internet Explorer browsers in Windows Vista or newer operating systems have an additional security feature called *Protected Mode*.

To launch and run ActiveX applications in Internet Explorer browsers with Protected Mode:

- 1. Run IE as an administrator.
- 2. Go to Tools > Internet Options > Security > Trusted Sites.
- Make sure that the Enable Protected Mode option is not selected for Trusted Sites zone. Alternatively, you can add the iDRAC address to sites in the Intranet zone. By default, protected mode is turned off for sites in Intranet Zone and Trusted Sites zone.
- 4. Click Sites.
- 5. In the Add this website to the zone field, add the address of your iDRAC and click Add.
- 6. Click **Close** and then click **OK**.
- 7. Close and restart the browser for the settings to take effect.

Clearing browser cache

If you have issues when operating the Virtual Console, (out of range errors, synchronization issues, and so on) clear the browser's cache to remove or delete any old versions of the viewer that may be stored on the system and try again.

Clearing earlier Java versions

To clear older versions of Java viewer in Windows or Linux, do the following:

- At the command prompt, run javaws-viewer or javaws-uninstall. The Java Cache viewer is displayed.
- 2. Delete the items titled iDRAC Virtual Console Client.

Importing CA certificates to management station

When you launch Virtual Console or Virtual Media, prompts are displayed to verify the certificates. If you have custom Web server certificates, you can avoid these prompts by importing the CA certificates to the Java or ActiveX trusted certificate store.

For more information about Automatic Certificate Enrollment (ACE), see section Automatic Certificate Enrollment

Importing CA certificate to Java trusted certificate store

To import the CA certificate to the Java trusted certificate store:

- 1. Launch the Java Control Panel.
- 2. Click Security tab and then click Certificates.
- The **Certificates** dialog box is displayed.
- 3. From the Certificate type drop-down menu, select **Trusted Certificates**.
- **4.** Click **Import**, browse, select the CA certificate (in Base64 encoded format), and click **Open**. The selected certificate is imported to the Web start trusted certificate store.
- 5. Click Close and then click OK. The Java Control Panel window closes.



Importing CA certificate to ActiveX trusted certificate store

You must use the OpenSSL command line tool to create the certificate Hash using Secure Hash Algorithm (SHA). It is recommended to use OpenSSL tool 1.0.x and later since it uses SHA by default. The CA certificate must be in Base64 encoded PEM format. This is one-time process to import each CA certificate.

To import the CA certificate to the ActiveX trusted certificate store:

- 1. Open the OpenSSL command prompt.
- 2. Run a 8 byte hash on the CA certificate that is currently in-use on the management station using the command: openssl x509 -in (name of CA cert) -noout -hash

An output file is generated. For example, if the CA certificate file name is **cacert.pem**, the command is:

openssl x509 -in cacert.pem -noout -hash

The output similar to "431db322" is generated.

- 3. Rename the CA file to the output file name and include a ".0" extension. For example, 431db322.0.
- 4. Copy the renamed CA certificate to your home directory. For example, C:\Documents and Settings\<user> directory.

Viewing localized versions of web interface

iDRAC web interface is supported in the following languages:

- English (en-us)
- French (fr)
- German (de)
- Spanish (es)
- Japanese (ja)
- Simplified Chinese (zh-cn)

The ISO identifiers in parentheses denote the supported language variants. For some supported languages, resizing the browser window to 1024 pixels wide is required to view all features.

iDRAC Web interface is designed to work with localized keyboards for the supported language variants. Some features of iDRAC Web interface, such as Virtual Console, may require additional steps to access certain functions or letters. Other keyboards are not supported and may cause unexpected problems.

NOTE: See the browser documentation on how to configure or setup different languages and view localized versions of iDRAC Web interface.

Updating device firmware

Using iDRAC, you can update the iDRAC, BIOS, and all device firmware that is supported by using Lifecycle Controller update such as:

- Fibre Channel (FC) cards
- Diagnostics
- Operating System Driver Pack
- Network Interface Card (NIC)
- RAID Controller
- Power Supply Unit (PSU)
- NVMe PCIe devices
- SAS/SATA hard drives
- Backplane update for internal and external enclosures
- OS Collector

CAUTION: The PSU firmware update may take several minutes depending on the system configuration and PSU model. To avoid damaging the PSU, do not interrupt the update process or power on the system during PSU firmware update.



NOTE: When updating the PSU firmware for PowerEdge C series servers, ensure that all servers in the same chassis are powered OFF first. If any of the other servers in the chassis are powered ON, the update process fails.

You must upload the required firmware to iDRAC. After the upload is complete, the current version of the firmware installed on the device and the version being applied is displayed. If the firmware being uploaded is not valid, an error message is displayed. Updates that do not require a reboot are applied immediately. Updates that require a system reboot are staged and committed to run on the next system reboot. Only one system reboot is required to perform all updates.

() NOTE:

- When SEKM mode is enabled on a controller, iDRAC Firmware downgrade/upgrade shall fail when tried from a SEKM to a non-SEKM iDRAC version. iDRAC Firmware upgrade/downgrade shall pass when done within the SEKM versions.
- PERC firmware downgrade shall fail when SEKM is enabled.

After the firmware is updated, the **System Inventory** page displays the updated firmware version and logs are recorded.

The supported firmware image file types are:

- .exe Windows-based Dell Update Package (DUP). You must have Control and Configure Privilege to use this image file type.
- .d9 Contains both iDRAC and Lifecycle Controller firmware

For files with .exe extension, you must have the System Control privilege. The Remote Firmware Update licensed feature and Lifecycle Controller must be enabled.

For files with .d9 extension, you must have the Configure privilege.

(i) NOTE: Ensure that all nodes in the system are powered off before updating the PSU firmware.

(i) **NOTE:** After upgrading the iDRAC firmware, you may notice a difference in the time stamp displayed in the Lifecycle Controller log. Time displayed in LC Log is different from NTP/Bios-Time for few logs during idrac reset.

You can perform firmware updates by using the following methods:

- Uploading a supported image type, one at a time, from a local system or network share.
- Connecting to an FTP, TFTP, HTTP or HTTPS site or a network repository that contains Windows DUPs and a corresponding catalog file.

You can create custom repositories by using the Dell Repository Manager. For more information, see *Dell Repository Manager Data Center User's Guide*. iDRAC can provide a difference report between the BIOS and firmware installed on the system and the updates available in the repository. All applicable updates contained in the repository are applied to the system. This feature is available with iDRAC Enterprise or Datacenter license.

(i) NOTE: HTTP/HTTPS only supports with either digest authentication or no authentication.

• Scheduling recurring automated firmware updates by using the catalog file and custom repository.

There are multiple tools and interfaces that can be used to update the iDRAC firmware. The following table is applicable only to iDRAC firmware. The table lists the supported interfaces, image-file types, and whether Lifecycle Controller must be in enabled state for the firmware to be updated.

Table 11. Image file types and dependencies

	.D9	.D9 Image		RAC DUP
Interface	Supported	Requires LC enabled	Supported	Requires LC enabled
BMCFW64.exe utility	Yes	No	No	N/A
Racadm FWUpdate (old)	Yes	No	No	N/A
Racadm Update (new)	Yes	Yes	Yes	Yes
iDRAC UI	Yes	Yes	Yes	Yes
WSMan	Yes	Yes	Yes	Yes
In-band OS DUP	No	N/A	Yes	No
Redfish	Yes	N/A	Yes	N/A



The following table provides information on whether a system restart is required when firmware is updated for a particular component:

() NOTE: When multiple firmware updates are applied through out-of-band methods, the updates are ordered in the most efficient possible manner to reduce unnecessary system restart.

Table 12. Firmware update — supported components

Component Name	Firmware Rollback Supported? (Yes or No)	Out-of-band — System Restart Required?	In-band — System Restart Required?	Lifecycle Controller GUI — Restart Required?
Diagnostics	No	No	No	No
OS Driver Pack	No	No	No	No
idrac	Yes	No	No*	Yes
BIOS	Yes	Yes	Yes	Yes
RAID Controller	Yes	Yes	Yes	Yes
BOSS	Yes	Yes	Yes	Yes
NVDIMM	No	Yes	Yes	Yes
Backplanes	Yes	Yes	Yes	Yes
 i) NOTE: • For Expander (Act • For SEP (Passive) 	ive) backplanes, system re backplanes, rebootless up	estart is required. date is supported only fror	n 4.00.00.00 release onwa	rds.
Enclosures	Yes	Yes	No	Yes
NIC	Yes	Yes	Yes	Yes
Power Supply Unit	Yes	Yes	Yes	Yes
CPLD	No	Yes	Yes	Yes
(i) NOTE: After CPLD fir	mware upgrade is complet	e, iDRAC restarts automat	tically.	
FC Cards	Yes	Yes	Yes	Yes
NVMe PCIe SSD drives	Yes	Yes	Yes	Yes
(i) NOTE: Rebootless up	date is supported on some	e devices starting from rele	ease 5.00.00.00.	
SAS/SATA hard drives	No	Yes	Yes	No
OS Collector	No	No	No	No
CMC (on PowerEdge FX2 servers)	No	Yes	Yes	Yes
ТРМ	No	Yes	Yes	Yes
(i) NOTE: TPM is suppor Downgrading and rein	ted from 5.00.00.00 relea stalling the same firmware	se onwards and the action is not supported.	is staged. Only firmware u	Ipdate is supported.

(i) NOTE: For details of supported components for MX platform, see Table 13.

Table 13. Firmware update — supported components for MX platforms

Component Name	Firmware Rollback Supported? (Yes or No)	Out-of-band — System Restart Required?	In-band — System Restart Required?	Lifecycle Controller GUI — Restart Required?
Diagnostics	No	No	No	No
OS Driver Pack	No	No	No	No

78 Setting up managed system



Table 13. Firmware update — supported components for MX platforms (continued)

Component Name	Firmware Rollback Supported? (Yes or No)	Out-of-band — System Restart Required?	In-band — System Restart Required?	Lifecycle Controller GUI — Restart Required?
idrac	Yes	No	No*	Yes
BIOS	Yes	Yes	Yes	Yes
RAID Controller	Yes	Yes	Yes	Yes
BOSS	Yes	Yes	Yes	Yes
NVDIMM	No	Yes	Yes	Yes
Backplanes	Yes	Yes	Yes	Yes
Enclosures	Yes	Yes	No	Yes
NIC	Yes	Yes	Yes	Yes
Power Supply Unit	No	No	No	No
CPLD	No	Yes	Yes	Yes
FC Cards	Yes	Yes	Yes	Yes
NVMe PCIe SSD drives	Yes	No	No	No
SAS/SATA hard drives	No	Yes	Yes	No
OS Collector	No	No	No	No

* Indicates that though a system restart is not required, iDRAC must be restarted to apply the updates. iDRAC communication and monitoring may temporarily be interrupted.

When you check for updates, the version marked as **Available** does not always indicate that it is the latest version available. Before you install the update, ensure that the version you choose to install is newer than the version currently installed. If you want to control the version that iDRAC detects, create a custom repository using Dell Repository Manager (DRM) and configure iDRAC to use that repository to check for updates.

Updating firmware using iDRAC web interface

You can update the device firmware using firmware images available on the local system, from a repository on a network share (CIFS, NFS, HTTP or HTTPS), or from FTP.

Updating single device firmware

Before updating the firmware using single device update method, make sure that you have downloaded the firmware image to a location on the local system.

(i) NOTE: Ensure that the file name for the single component DUP does not have any blank space.

To update single device firmware using iDRAC web interface:

1. Go to Maintenance > System Update.

The Firmware Update page is displayed.

2. On the Update tab, select Local as the Location Type.

NOTE: If you select Local, ensure that you download the firmware image to a location on the local system. Select one file to be staged to iDRAC for update. You can select additional files one file at a time, for uploading to iDRAC. The files are uploaded to a temporary space on iDRAC and is limited to approximately 300MB.

- 3. Click Browse, select the firmware image file for the required component, and then click Upload.
- 4. After the upload is complete, the Update Details section displays each firmware file uploaded to iDRAC and its status.



If the firmware image file is valid and was successfully uploaded, the **Contents** column displays a plus icon (+) icon next to the firmware image file name. Expand the name to view the **Device Name**, **Current**, and **Available firmware version** information.

- 5. Select the required firmware file and do one of the following:
 - For firmware images that do not require a host system reboot, click **Install** (only available option). For example, iDRAC firmware file.
 - For firmware images that require a host system reboot, click Install and Reboot or Install Next Reboot.
 - To cancel the firmware update, click **Cancel**.

When you click **Install and Reboot**, or **Install Next Reboot**, the message Updating Job Queue is displayed.

To display the Job Queue page, click Job Queue. Use this page to view and manage the staged firmware updates or click OK to refresh the current page and view the status of the firmware update.

NOTE: If you navigate away from the page without saving the updates, an error message is displayed and all the uploaded content is lost.

(i) **NOTE:** You will not be able to proceed further, if the session gets expired after uploading the firmware file. This issue can be only resolved by RACADM reset.

() NOTE: After firmaware update is completed, an error message is displayed - RAC0508: An unexpected error

occurred. Wait for few minutes and retry the operation. If the problem persists, contact service provider. This is expected. You can wait for sometime and refresh the browser. Then you are re-directed to login page.

Scheduling automatic firmware updates

You can create a periodic recurring schedule for iDRAC to check for new firmware updates. At the scheduled date and time, iDRAC connects to the specified destination, checks for new updates, and applies or stages all applicable updates. A log file is created on the remote server, which contains information about server access and staged firmware updates.

It is recommended that you create a repository using Dell Repository Manager (DRM) and configure iDRAC to use this repository to check for and perform firmware updates. Using an internal repository enables you to control the firmware and versions available to iDRAC and helps avoid any unintended firmware changes.

(i) NOTE: For more information about DRM, see www.dell.com/openmanagemanuals > Repository Manager .

iDRAC Enterprise or Datacenter license is required to schedule automatic updates.

You can schedule automatic firmware updates using the iDRAC web interface or RACADM.

(i) NOTE: IPv6 address is not supported for scheduling automatic firmware updates.

Scheduling automatic firmware update using web interface

To schedule automatic firmware update using web Interface:

- () NOTE: Do not create the next scheduled occurrence of an automatic update job if a job is already Scheduled. It overwrites the current scheduled job.
- In the iDRAC web interface, go to Maintenance > System Update > Automatic Update. The Firmware Update page is displayed.
- 2. Click the Automatic Update tab.
- 3. Select the Enable Automatic Update option.
- 4. Select any of the following options to specify if a system reboot is required after the updates are staged:
 - Schedule Updates Stage the firmware updates but do not reboot the server.
 - Schedule Updates and reboot Server Enables server reboot after the firmware updates are staged.
- 5. Select any of the following to specify the location of the firmware images:
 - **Network** Use the catalog file from a network share (CIFS, NFS, HTTP or HTTPS, TFTP). Enter the network share location details.

NOTE: While specifying the network share settings, it is recommended to avoid special characters for user name and password or percent encode the special characters.



- **FTP** Use the catalog file from the FTP site. Enter the FTP site details.
- **HTTP** or **HTTPS** Allows catalog file streaming and via HTTP and HTTPS file transfer.

6. Based on the selection in step 5, enter the network settings or the FTP settings. For information about the fields, see the *iDRAC Online Help*.

7. In the **Update Window Schedule** section, specify the start time for the firmware update and the frequency of the updates (daily, weekly, or monthly).

For information about the fields, see the *iDRAC Online Help*.

8. Click Schedule Update.

The next scheduled job is created in the job queue. Five minutes after the first instance of the recurring job starts, the job for the next time period is created.

Scheduling automatic firmware update using RACADM

To schedule automatic firmware update, use the following commands:

• To enable automatic firmware update:

racadm set lifecycleController.lcattributes.AutoUpdate.Enable 1

• To view the status of automatic firmware update:

racadm get lifecycleController.lcattributes.AutoUpdate

• To schedule the start time and frequency of the firmware update:

```
racadm AutoUpdateScheduler create -u username -p password -l <location> [-f
catalogfilename -pu <proxyuser> -pp<proxypassword> -po <proxy port> -pt <proxytype>]
-time < hh:mm> [-dom < 1 - 28,L,'*'> -wom <l-4,L,'*'> -dow <sun-sat,'*'>] -rp <l-366>
-a <applyserverReboot (1-enabled | 0-disabled)>
```

For example,

• To automatically update firmware using a CIFS share:

racadm AutoUpdateScheduler create -u admin -p pwd -l //1.2.3.4/CIFS-share -f cat.xml -time 14:30 -wom 1 -dow sun -rp 5 -a 1

To automatically update firmware using FTP:

racadm AutoUpdateScheduler create -u admin -p pwd -l ftp.mytest.com -pu puser -pp puser -po 8080 -pt http -f cat.xml -time 14:30 -wom 1 -dow sun -rp 5 -a 1

• To view the current firmware update schedule:

racadm AutoUpdateScheduler view

• To disable automatic firmware update:

racadm set lifecycleController.lcattributes.AutoUpdate.Enable 0

• To clear the schedule details:

racadm AutoUpdateScheduler clear

Updating device firmware using RACADM

To update device firmware using RACADM, use the update subcommand. For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Examples:

• Upload the update file from a remote HTTP share:

```
racadm update -f <updatefile> -u admin -p mypass -l http://1.2.3.4/share
```



• Upload the update file from a remote HTTPS share:

racadm update -f <updatefile> -u admin -p mypass -l https://1.2.3.4/share

• To generate a comparison report using an update repository:

```
racadm update -f catalog.xml -l //192.168.1.1 -u test -p passwd --verifycatalog
```

 To perform all applicable updates from an update repository using myfile.xml as a catalog file and perform a graceful reboot:

racadm update -f "myfile.xml" -b "graceful" -l //192.168.1.1 -u test -p passwd

To perform all applicable updates from an FTP update repository using Catalog.xml as a catalog file:

racadm update -f "Catalog.xml" -t FTP -e 192.168.1.20/Repository/Catalog

Updating firmware using CMC web interface

You can update iDRAC firmware for blade servers using the CMC Web interface.

To update iDRAC firmware using the CMC Web interface:

- 1. Log in to CMC Web interface.
- Go to iDRAC Settings > Settings > CMC. The Deploy iDRAC page is displayed.
- 3. Click Launch iDRAC Web interface and perform iDRAC Firmware Update.

Updating firmware using DUP

Before you update firmware using Dell Update Package (DUP), make sure to:

- Install and enable the IPMI and managed system drivers.
- Enable and start the Windows Management Instrumentation (WMI) service if your system is running Windows operating system,
 - **NOTE:** While updating the iDRAC firmware using the DUP utility in Linux, if you see error messages such as usb 5-2:
 - device descriptor read/64, error -71 displayed on the console, ignore them.
- If the system has ESX hypervisor installed, then for the DUP file to run, make sure that the "usbarbitrator" service is stopped using command: service usbarbitrator stop

Some versions of DUPs are constructed in ways that conflict with each other. This happens over time as new versions of the software are created. A newer version of software may drop support for legacy devices. Support for new devices may be added. Consider, for example, the two DUPs Network_Firmware_NDT09_WN64_21.60.5.EXE and Network_Firmware_8J1P7_WN64_21.60.27.50.EXE. The devices supported by these DUPs fit into three groups.

- Group A are legacy devices supported only by NDT09.
- Group B are devices supported by both NDT09 and 8J1P7.
- Group C are new devices supported only by 8J1P7.

Consider a server that has one or more devices from each of Groups A, B, and C. If the DUPs are used one at a time they should be successful. Using NDT09 by itself updates the devices in group A and group B. Using 8J1P7 by itself updates devices in group B and group C. However, if you try to use both DUPs at the same time that may attempt to create two updates for the Group B devices at the same time. That may fail with a valid error: "Job for this device is already present". The update software is unable to resolve the conflict of two valid DUPs attempting two valid updates on the same devices at the same time. At the same time both DUPs are required to support Group A and Group C devices. The conflict extends to performing rollbacks on the devices too. For best practice it is suggested to use each DUP individually.

To update iDRAC using DUP:

- 1. Download the DUP based on the installed operating system and run it on the managed system.
- 2. Run the DUP.
- The firmware is updated. A system restart is not required after firmware update is complete.



Updating firmware using remote RACADM

- 1. Download the firmware image to the TFTP or FTP server. For example, C:\downloads\firmimg.d9
- 2. Run the following RACADM command:

```
TFTP server:
```

• Using fwupdate command:

```
racadm -r <iDRAC IP address> -u <username> -p <password> fwupdate -g -u -a <path>
```

path

the location on the TFTP server where firmimg.d9 is stored.

• Using update command:

```
racadm -r <iDRAC IP address> -u <username> -p <password> update -f <filename>
```

FTP server:

• Using fwupdate command:

```
racadm -r <iDRAC IP address> -u <username> -p <password> fwupdate -f <ftpserver IP>
<ftpserver username> <ftpserver password> -d <path>
```

path

the location on the FTP server where firmimg.d9 is stored.

Using update command:

racadm -r <iDRAC IP address> -u <username> -p <password> update -f <filename>

For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Updating firmware using Lifecycle Controller Remote Services

For information to update the firmware using Lifecycle Controller–Remote Services, see *Lifecycle Controller Remote Services Quick Start Guide* available at https://www.dell.com/idracmanuals.

Updating CMC firmware from iDRAC

In the PowerEdge FX2/FX2s chassis, you can update the firmware for the Chassis Management Controller and any component that can be updated by CMC and shared by the servers from iDRAC.

Before applying the update, make sure that:

- Servers are not allowed to power-up by CMC.
- Chassis with LCD must display a message indicating "update is in-progress".
- Chassis without LCD must indicate the update progress using LED blinking pattern.
- During the update, chassis action power commands are disabled.

The updates for components such as Programmable System-on-Chip (PSoC) of IOM that requires all the servers to be idle, the update is applied on the next chassis power-up cycle.

CMC settings to update CMC firmware from iDRAC

In the PowerEdge FX2/FX2s chassis, before performing the firmware update from iDRAC for CMC and its shared components, do the following:

- 1. Launch the CMC Web interface
- Go to iDRAC Settings > Settings > CMC. The Deploy iDRAC page is displayed.
- 3. From the Chassis Management at Server Mode , select Manage and Monitor, and the click Apply.



iDRAC settings to update CMC firmware

In the PowerEdge FX2/FX2s chassis, before updating the firmware for CMC and its shared components from iDRAC, do the following settings in iDRAC:

- 1. Go to iDRAC Settings > Settings > CMC.
- 2. Click on Chassis Management Controller Firmware Update The Chassis Management Controller Firmware Update Settings page is displayed.
- For Allow CMC Updates Through OS and Lifecycle Controller, select Enabled to enable CMC firmware update from iDRAC.
- 4. Under Current CMC Setting, make sure that Chassis Management at Server Mode option displays Manage and Monitor. You can set this in CMC.

Viewing and managing staged updates

You can view and delete the scheduled jobs including configuration and update jobs. This is a licensed feature. All jobs queued to run during the next reboot can be deleted.

Viewing and managing staged updates using iDRAC web interface

To view the list of scheduled jobs using iDRAC web interface, go to **Maintenance** > **Job Queue**. The **Job Queue** page displays the status of jobs in the Lifecycle Controller job queue. For information about the displayed fields, see the *iDRAC Online Help*.

To delete job(s), select the job(s) and click **Delete**. The page is refreshed and the selected job is removed from the Lifecycle Controller job queue. You can delete all the jobs queued to run during the next reboot. You cannot delete active jobs, that is, jobs with the status *Running* or *Downloading*.

You must have Server Control privilege to delete jobs.

Viewing and managing staged updates using RACADM

To view the staged updates using RACADM, use jobqueue sub-command. For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Rolling back device firmware

You can roll back the firmware for iDRAC or any device that Lifecycle Controller supports, even if the upgrade was previously performed using another interface. For example, if the firmware was upgraded using the Lifecycle Controller GUI, you can roll back the firmware using the iDRAC web interface. You can perform firmware rollback for multiple devices with one system reboot.

On Dell's 14th generation PowerEdge servers that have a single iDRAC and Lifecycle Controller firmware, rolling back the iDRAC firmware also rolls back the Lifecycle Controller firmware.

It is recommended to keep the firmware updated to ensure you have the latest features and security updates. You may need to rollback an update or install an earlier version if you encounter any issues after an update. To install an earlier version, use Lifecycle Controller to check for updates and select the version you want to install.

For the details about supported and unsupported components for firmware rollback, refer to the table Firmware update — supported components

You can perform firmware rollback for the following components:

- iDRAC with Lifecycle Controller
- BIOS
- Network Interface Card (NIC)
- Power Supply Unit (PSU)
- RAID Controller
- Backplane



(i) NOTE: You cannot perform firmware rollback for Diagnostics, Driver Packs, and CPLD.

Before rolling back the firmware, make sure that:

- You have Configure privilege to roll back iDRAC firmware.
- You have Server Control privilege and have enabled Lifecycle Controller to roll back firmware for any other device other than the iDRAC.
- Change the NIC mode to **Dedicated** if the mode is set as **Shared LOM**.

You can roll back the firmware to the previously installed version using any of the following methods:

- iDRAC web interface
- CMC web interface (not supported on MX platforms)
- OME-Modular web interface (Supported on MX platforms)
- CMC RACADM CLI (not supported on MX platforms)
- iDRAC RACADM CLI
- Lifecycle Controller GUI
- Lifecycle Controller-Remote Services

Rollback firmware using iDRAC web interface

To roll back device firmware:

1. In the iDRAC Web interface, go to Maintenance > System Update > Rollback.

The **Rollback** page displays the devices for which you can rollback the firmware. You can view the device name, associated devices, currently installed firmware version, and the available firmware rollback version.

- 2. Select one or more devices for which you want to rollback the firmware.
- 3. Based on the selected devices, click Install and Reboot or Install Next Reboot. If only iDRAC is selected, then click Install.
 - When you click Install and Reboot or Install Next Reboot, the message "Updating Job Queue" is displayed.

4. Click Job Queue.

The **Job Gueue** page is displayed, where you can view and manage the staged firmware updates.

() NOTE:

• While in rollback mode, the rollback process continues in the background even if you navigate away from this page.

An error message appears if:

- You do not have Server Control privilege to rollback any firmware other than the iDRAC or Configure privilege to rollback iDRAC firmware.
- Firmware rollback is already in-progress in another session.
- Updates are staged to run or already in running state.

If Lifecycle Controller is disabled or in recovery state and you try to perform a firmware rollback for any device other than iDRAC, an appropriate warning message is displayed along with steps to enable Lifecycle Controller.

Rollback firmware using CMC web interface

To roll back using the CMC Web interface:

- 1. Log in to CMC Web interface.
- Go to iDRAC Settings > Settings > CMC. The Deploy iDRAC page is displayed.
- **3.** Click **Launch iDRAC** and perform device firmware rollback as mentioned in the Rollback firmware using iDRAC web interface.

Rollback firmware using RACADM

1. Check the rollback status and the FQDD using the swinventory command:

racadm swinventory



For the device for which you want to rollback the firmware, the Rollback Version must be Available. Also, note the FQDD.

2. Rollback the device firmware using:

racadm rollback <FQDD>

For more information, see *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Rollback firmware using Lifecycle Controller

For information, see Lifecycle Controller User's Guide available at https://www.dell.com/idracmanuals .

Rollback firmware using Lifecycle Controller-Remote Services

For information, see Lifecycle Controller Remote Services Quick Start Guide available at https://www.dell.com/idracmanuals.

Recovering iDRAC

iDRAC supports two operating system images to make sure a bootable iDRAC. In the event of an unforeseen catastrophic error and you lose both boot paths:

- iDRAC bootloader detects that there is no bootable image.
- System Health and Identify LED is flashed at ~1/2 second rate. (LED is located on the back of a rack and tower servers and on the front of a blade server.)
- Bootloader is now polling the SD card slot.
- Format an SD card with FAT using a Windows operating system, or EXT3 using a Linux operating system.
- Copy **firmimg.d9** to the SD card.
- Insert the SD card into the server.
- Bootloader detects the SD card, turns the flashing LED to solid amber, reads the firmimg.d9, reprograms iDRAC, and then reboots iDRAC.

Monitoring iDRAC using other Systems Management tools

You can discover and monitor iDRAC using Dell Management Console or Dell OpenManage Essentials. You can also use Dell Remote Access Configuration Tool (DRACT) to discover iDRACs, update firmware, and set up Active Directory. For more information, see the respective user's guides.

Support Server Configuration Profile — Import and Export

Server Configuration Profile (SCP) allows you to import and export server configuration files.

(i) NOTE: You need admin privileges to perform Export and Import SCP task.

You can import and export from local management station, and from a Network Share via CIFS, NFS, HTTP or HTTPS. Using SCP, you can select and import or export component level configurations for BIOS, NIC and RAID. You can import and export SCP to the local management station or to a CIFS, NFS, HTTP, or HTTPS network share. You can either import and export individual profiles of iDRAC, BIOS, NIC, and RAID, or all of them together as a single file.

You can specify preview import or export of the SCP where the job is running and configuration result is generated but none of the configuration has applied.

A job is created once the import or export is initiated through the GUI. The status of the jobs can be viewed on the Job Queue page.

86 Setting up managed system



(i) NOTE: Only Host Name or IP Address are accepted for destination address.

- **NOTE:** You can browse to a specific location to import the server configuration files. You need to select the correct server configuration file that you want to import. For example, import.xml.
- **NOTE:** Depending on the exported file format (that you selected), the extension is added automatically. For example, export_system_config.xml.
- () NOTE: SCP applies the full configuration in a single job with minimal number of reboots. However, in a few system configurations some attributes change the operation mode of a device or may create subdevices with new attributes. When this occurs, SCP may be unable to apply all settings during a single job. Review the ConfigResult entries for the job to resolve any pending configuration settings.

SCP allows you to perform OS deployment (OSD) using a single xml/json file across multiple systems. You can also perform existing operations such as configurations and repository updates all at once.

SCP also allows to export and import of SSH public keys for all iDRAC users. There are 4 SSH public keys for all users.

Following are the steps for OS deployment using SCP:

- 1. Export SCP file
- 2. SCP file contains all the suppressed attributes that are needed to perform OSD.
- 3. Edit / update the OSD attributes and then perform import operation.
- 4. These OSD attributes are then validated by SCP orchestrator.
- 5. SCP orchestrator performs the configuration and repository updates specified in SCP file.
- 6. After configuration and updates are done, host OS shutdowns.

(i) NOTE: Only CIFS and NFS share is supported for hosting OS media.

- 7. SCP orchestrator initiates the OSD by attaching the drivers for the selected operating system and then initiates one time boot to the OS media present in NFS / Share.
- $\textbf{8.} \hspace{0.1 cm} \text{LCL shows the progress of the job.}$
- ${\bf 9.}~$ Once BIOS boots to the OS media, SCP job shows as Complete.
- 10. The attached media and OS media will be automatically detached after 65535 seconds or after the duration specified by OSD.1#ExposeDuration attribute.

Importing server configuration profile using iDRAC web interface

To import the server configuration profile:

- 1. Go to Configuration > Server Configuration Profile The Server Configuration Profile page is displayed.
- 2. Select one of the following to specify the location type:
 - Local to import the configuration file saved in a local drive.
 - **Network Share** to import the configuration file from CIFS or NFS share.
 - HTTP or HTTPS to import the configuration file from a local file using HTTP/HTTPS file transfer.
 - **NOTE:** Depending on the location type, you must enter the Network Settings or HTTP/HTTPS settings. If proxy is configured for HTTP/HTTPS, proxy settings are also required.
- 3. Select the components listed in Import Components option.
- 4. Select the **Shutdown** type.
- 5. Select the **Maximum wait time** to specify the wait time before the system shuts down after the import is complete.
- 6. Click Import.

Exporting server configuration profile using iDRAC web interface

To export the server configuration profile:

- 1. Go to Configuration > Server Configuration Profile The Server Configuration Profile page is displayed.
- 2. Click Export.



- 3. Select one of the following to specify the location type:
 - Local to save the configuration file on a local drive.
 - **Network Share** to save the configuration file on a CIFS or NFS share.
 - **HTTP or HTTPS** to save the configuration file to a local file using HTTP/HTTPS file transfer.
 - () NOTE: Depending on the location type, you must enter the Network Settings or HTTP/HTTPS settings. If proxy is configured for HTTP/HTTPS, proxy settings are also required.
- 4. Select the components that you need to back up the configuration for.
- 5. Select the **Export type**, following are the options:
 - Basic
 - Replacement Export
 - Clone Export
- 6. Select an Export file format.
- 7. Select Additional export items.
- 8. Click Export.

Secure Boot Configuration from BIOS Settings or F2

UEFI Secure Boot is a technology that eliminates a major security void that may occur during a handoff between the UEFI firmware and UEFI operating system (OS). In UEFI Secure Boot, each component in the chain is validated and authorized against a specific certificate before it is allowed to load or run. Secure Boot removes the threat and provides software identity checking at every step of the boot—Platform firmware, Option Cards, and OS BootLoader.

The Unified Extensible Firmware Interface (UEFI) Forum—an industry body that develops standards for pre-boot software —defines Secure Boot in the UEFI specification. Computer system vendors, expansion card vendors, and operating system providers collaborate on this specification to promote interoperability. As a portion of the UEFI specification, Secure Boot represents an industry-wide standard for security in the pre-boot environment.

When enabled, UEFI Secure Boot prevents the unsigned UEFI device drivers from being loaded, displays an error message, and does not allow the device to function. You must disable Secure Boot to load the unsigned device drivers.

On the Dell 14th generation and later versions of PowerEdge servers, you can enable or disable the Secure Boot feature by using different interfaces (RACADM, WSMAN, REDFISH, and LC-UI).

Acceptable file formats

The Secure Boot policy contains only one key in PK, but multiple keys may reside in KEK. Ideally, either the platform manufacturer or platform owner maintains the private key corresponding to the public PK. Third parties (such as OS providers and device providers) maintain the private keys corresponding to the public keys in KEK. In this way, platform owners or third parties may add or remove entries in the db or dbx of a specific system.

The Secure Boot policy uses db and dbx to authorize pre-boot image file execution. For an image file to get executed, it must associate with a key or hash value in db, and not associate with a key or hash value in dbx. Any attempts to update the contents of db or dbx must be signed by a private PK or KEK. Any attempts to update the contents of PK or KEK must be signed by a private PK.

Table 14. Acceptable file formats

Policy Component	Acceptable File Formats	Aco	ceptable File Extensions	Max records allowed
РК	X.509 Certificate (binary DER format only)	1.	.cer	One
		2.	.der	
		3.	.crt	
КЕК	X.509 Certificate (binary DER format only)	1.	.cer	More than one
	Public Key Store	2.	.der	
		3.	.crt	



Table 14. Acceptable file formats (continued)

Policy Component	Acceptable File Formats	Ace	ceptable File Extensions	Max records allowed
		4.	.pbk	
DB and DBX	X.509 Certificate (binary DER format only)	1.	.cer	More than one
	EFI image (system BIOS will calculate and import image digest)	2.	.der	
		3.	.crt	
		4.	.efi	

The Secure Boot Settings feature can be accessed by clicking System Security under System BIOS Settings. To go to System BIOS Settings, press F2 when the company logo is displayed during POST.

- By default, Secure Boot is Disabled and the Secure Boot policy is set to Standard. To configure the Secure Boot Policy, you must enable Secure Boot.
- When the Secure Boot mode is set to Standard, it indicates that the system has default certificates and image digests or hash loaded from the factory. This caters to the security of standard firmware, drivers, option-roms, and boot loaders.
- To support a new driver or firmware on a server, the respective certificate must be enrolled into the DB of Secure Boot certificate store. Therefore, Secure Boot Policy must be configured to Custom.

When the Secure Boot Policy is configured as Custom, it inherits the standard certificates and image digests loaded in the system by default, which you can modify. Secure Boot Policy configured as Custom allows you to perform operations such as View, Export, Import, Delete, Delete All, Reset, and Reset All. Using these operations, you can configure the Secure Boot Policies.

Configuring the Secure Boot Policy to Custom enables the options to manage the certificate store by using various actions such as Export, Import, Delete, Delete All, Reset, and Rest All on PK, KEK, DB, and DBX. You can select the policy (PK / KEK / DB / DBX) on which you want to make the change and perform appropriate actions by clicking the respective link. Each section will have links to perform the Import, Export, Delete, and Reset operations. Links are enabled based on what is applicable, which depends on the configuration at the point of time. Delete All and Reset All are the operations that have impact on all the policies. Delete All deletes all the certificates and image digests in the Custom policy, and Rest All restores all the certificates and image digests from Standard or Default certificate store.

BIOS recovery

The BIOS recovery feature allows you to manually recover the BIOS from a stored image. The BIOS is checked when the system is powered-on and if a corrupt or compromised BIOS is detected, an error message is displayed. You can then initiate the process of BIOS recovery using RACADM. To perform a manual BIOS recovery, see the iDRAC RACADM Command Line Interface Reference Guide available at https://www.dell.com/idracmanuals.



Configuring iDRAC

iDRAC enables you to configure iDRAC properties, set up users, and set up alerts to perform remote management tasks.

Before you configure iDRAC, make sure that the iDRAC network settings and a supported browser is configured, and the required licenses are updated. For more information about the licensable feature in iDRAC, see iDRAC licenses .

You can configure iDRAC using:

- iDRAC Web Interface
- RACADM
- Remote Services (see Lifecycle Controller Remote Services User's Guide)
- IPMITool (see Baseboard Management Controller Management Utilities User's Guide)

To configure iDRAC:

- 1. Log in to iDRAC.
- **2.** Modify the network settings if required.

NOTE: If you have configured iDRAC network settings, using iDRAC Settings utility during iDRAC IP address setup, then ignore this step.

- 3. Configure interfaces to access iDRAC.
- 4. Configure front panel display.
- 5. Configure System Location if required.
- 6. Configure time zone and Network Time Protocol (NTP) if required.
- 7. Establish any of the following alternate communication methods to iDRAC:
 - IPMI or RAC serial
 - IPMI serial over LAN
 - IPMI over LAN
 - SSH
- 8. Obtain the required certificates.
- 9. Add and configure iDRAC users with privileges.
- 10. Configure and enable e-mail alerts, SNMP traps, or IPMI alerts.
- 11. Set the power cap policy if required.
- **12.** Enable the Last Crash Screen.
- 13. Configure virtual console and virtual media if required.
- 14. Configure vFlash SD card if required.
- **15.** Set the first boot device if required.
- **16.** Set the OS to iDRAC Pass-through if required.

Topics:

- Viewing iDRAC information
- Modifying network settings
- Cipher suite selection
- FIPS mode
- Configuring services
- Using VNC client to manage remote server
- Configuring front panel display
- Configuring time zone and NTP
- Setting first boot device
- Enabling or disabling OS to iDRAC Pass-through
- Obtaining certificates
- Configuring multiple iDRACs using RACADM



Disabling access to modify iDRAC configuration settings on host system

Viewing iDRAC information

You can view the basic properties of iDRAC.

Viewing iDRAC information using web interface

In the iDRAC Web interface, go to **iDRAC Settings** > **Overview** to view the following information related to iDRAC. For information about the properties, see *iDRAC Online Help*.

iDRAC Details

- Device Type
- Hardware Version
- Firmware Version
- Firmware Update
- RAC time
- IPMI version
- Number of Possible Sessions
- Number of Current Sessions
- IPMI Version

iDRAC Service Module

Status

Connection View

- State
- Switch Connection ID
- Switch Port Connection ID

Current Network Settings

- iDRAC MAC Address
- Active NIC Interface
- DNS Domain Name

Current IPv4 Setting

- IPv4 Enabled
- DHCP
- Current IP Address
- Current Subnet Mask
- Current Gateway
- Use DHCP to Obtain DNS Server Address
- Current Preferred DNS Server
- Current Alternate DNS Server

Current IPv6 Settings

- IPv6 Enable
- Autoconfiguration
- Current IP Address
- Current IP Gateway
- Link Local Address
- Use DHCPv6 to obtain DNS
- Current Preferred DNS Server
- Current Alternate DNS Server



Viewing iDRAC information using RACADM

To view iDRAC information using RACADM, see getsysinfo or get sub-command details provided in the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Modifying network settings

After configuring the iDRAC network settings using the iDRAC Settings utility, you can also modify the settings through the iDRAC Web interface, RACADM, Lifecycle Controller, and Server Administrator (after booting to the operating system). For more information on the tools and privilege settings, see the respective user's guides.

To modify the network settings using iDRAC Web interface or RACADM, you must have **Configure** privileges.

(i) NOTE: Changing the network settings may terminate the current network connections to iDRAC.

Modifying network settings using web interface

To modify the iDRAC network settings:

- In the iDRAC Web interface, go to iDRAC Settings > Connectivity > Network > Network Settings. The Network page is displayed.
- 2. Specify the network settings, common settings, IPv4, IPv6, IPMI, and/or VLAN settings as per your requirement and click **Apply**.

If you select **Auto Dedicated NIC** under **Network Settings**, when the iDRAC has its NIC Selection as shared LOM (1, 2, 3, or 4) and a link is detected on the iDRAC dedicated NIC, the iDRAC changes its NIC selection to use the dedicated NIC. If no link is detected on the dedicated NIC, then the iDRAC uses the shared LOM. The switch from shared to dedicated time-out is five seconds and from dedicated to shared is 30 seconds. You can configure this time-out value using RACADM or WSMan.

For information about the various fields, see the *iDRAC Online Help*.

NOTE: If the iDRAC is using DHCP and has obtained a lease for its IP address, it is released back to the DHCP server's address pool when NIC or Ipv4 or DHCP is disabled.

Modifying network settings using local RACADM

To generate a list of available network properties, use the command

racadm get iDRAC.Nic

To use DHCP to obtain an IP address, use the following command to write the object DHCPEnable and enable this feature.

racadm set iDRAC.IPv4.DHCPEnable 1

The following example shows how the command may be used to configure the required LAN network properties:

racadm	set	iDRAC.Nic.Enable 1
racadm	set	iDRAC.IPv4.Address 192.168.0.120
racadm	set	iDRAC.IPv4.Netmask 255.255.0
racadm	set	iDRAC.IPv4.Gateway 192.168.0.120
racadm	set	iDRAC.IPv4.DHCPEnable 0
racadm	set	iDRAC.IPv4.DNSFromDHCP 0
racadm	set	iDRAC.IPv4.DNS1 192.168.0.5
racadm	set	iDRAC.IPv4.DNS2 192.168.0.6
racadm	set	iDRAC.Nic.DNSRegister 1
racadm	set	iDRAC.Nic.DNSRacName RAC-EK00002
racadm	set	iDRAC.Nic.DNSDomainFromDHCP 0
racadm	set	iDRAC.Nic.DNSDomainName MYDOMAIN

(i) NOTE: If iDRAC.Nic.Enable is set to **0**, the iDRAC LAN is disabled even if DHCP is enabled.



Configuring IP filtering

In addition to user authentication, use the following options to provide additional security while accessing iDRAC:

- IP filtering limits the IP address range of the clients accessing iDRAC. It compares the IP address of an incoming login to
 the specified range and allows iDRAC access only from a management station whose IP address is within the range. All other
 login requests are denied.
- When repeated login failures occur from a particular IP address, it prevents the address from logging in to iDRAC for a preselected time span. If you unsuccessfully log in up to two times, you are allowed to log in again only after 30 seconds. If you unsuccessfully log in more than two times, you are allowed to log in again only after 60 seconds.

(i) NOTE: This features supports upto 5 IP ranges. You van view / set this feature using RACADM and Redfish.

As login failures accumulate from a specific IP address, they are registered by an internal counter. When the user successfully logs in, the failure history is cleared and the internal counter is reset.

NOTE: When login attempts are prevented from the client IP address, few SSH clients may display the message: ssh exchange identification: Connection closed by remote host.

Configure IP filtering using iDRAC web interface

You must have Configure privilege to perform these steps.

To configure IP filtering:

1. In iDRAC Web interface, go to iDRAC Settings > Connectivity > Network > Network Settings > Advanced Network Settings.

The **Network** page is displayed.

- Click Advanced Network Settings. The Network Security page is displayed.
- **3.** Specify the IP filtering settings using **IP Range Address** and **IP Range Subnet Mask**. For more information about the options, see *iDRAC Online Help*.
- 4. Click Apply to save the settings.

Federal Information Processing Standards — FIPS is a set of standards used by the United States government agencies and contractors. FIPS Mode is intended to meet the requirements of FIPS 140-2 level 1. For more information about FIPS, refer to the FIPS User Guide for iDRAC, and CMC for non MX platforms.

(i) NOTE: Enabling FIPS Mode resets iDRAC to the default settings.

Configuring IP filtering using RACADM

You must have Configure privilege to perform these steps.

To configure IP filtering, use the following RACADM objects in the iDRAC.IPBlocking group:

- RangeEnable
- RangeAddr
- RangeMask

The RangeMask property is applied to both the incoming IP address and to the RangeAddr property. If the results are identical, the incoming login request is allowed to access iDRAC. Logging in from IP addresses outside this range results in an error.

(i) NOTE: Configuring IP filtering supports up to 5 IP ranges.

The login proceeds if the following expression equals zero:

```
RangeMask & (<incoming-IP-address> ^ RangeAddr)
```

&

Bitwise AND of the quantities

^



Bitwise exclusive-OR

Examples for IP Filtering

The following RACADM commands block all IP addresses except 192.168.0.57:

```
racadm set iDRAC.IPBlocking.RangeEnable 1
racadm set iDRAC.IPBlocking.RangeAddr 192.168.0.57
racadm set iDRAC.IPBlocking.RangeMask 255.255.255.255
```

To restrict logins to a set of four adjacent IP addresses (for example, 192.168.0.212 through 192.168.0.215), select all but the lowest two bits in the mask:

racadm set iDRAC.IPBlocking.RangeEnable 1
racadm set iDRAC.IPBlocking.RangeAddr 192.168.0.212
racadm set iDRAC.IPBlocking.RangeMask 255.255.255.252

The last byte of the range mask is set to 252, the decimal equivalent of 11111100b.

For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Cipher suite selection

Cipher Suite Selection can be used to limit the ciphers in iDRAC or client communications and determine how secure the connection will be. It provides another level of filtering the effective in-use TLS Cipher Suite. These settings can be configured through iDRAC web interface, RACADM, and WSMan command line interfaces.

Configuring cipher suite selection using iDRAC web interface

CAUTION: Using OpenSSL Cipher Command to parse strings with invalid syntax may lead to unexpected errors.

NOTE: This is an advanced security option. Before you configure this option, ensure that you have thorough knowledge of the following:

- The OpenSSL Cipher String Syntax and its use.
- Tools and Procedures to validate the resultant Cipher Suite Configuration to ensure that the results align with the
 expectations and requirements.
- **NOTE:** Before you configure the Advanced Settings for TLS Cipher Suites, ensure that you are using a supported web browser.

To add custom cipher strings:

- 1. In iDRAC web interface, go to iDRAC Settings > Services > Web Server.
- 2. Click Set Cipher String under the Customer Cipher String option.
- The Set Custom Cipher String page is displayed.
- 3. In the Custom Cipher String field, enter a valid string and click Set Cipher String.

(i) NOTE: For more information about cipher strings, see www.openssl.org/docs/man1.0.2/man1/ciphers.html.

4. Click Apply.

Setting the custom cipher string terminates the current iDRAC session. Wait for a few minutes before you open a new iDRAC session.

The ciphers supported by iDRAC on port 5000 are:

ssl-enum-ciphers:

TSLv1.1 Ciphers:

- TLS_ECDHE_RSA_WITH_3DES_EDE_CBC_SHA (secp256r1)
- TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA (secp256r1)
- TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA (secp256r1)
- TLS_ECDHE_RSA_WITH_RC4_128_SHA (secp256r1)
- TLS_RSA_WITH_3DES_EDE_CBC_SHA (rsa 2048)

94 Configuring iDRAC



- TLS_RSA_WITH_AES_128_CBC_SHA (rsa 2048)
- TLS_RSA_WITH_AES_256_CBC_SHA (rsa 2048)
- TLS_RSA_WITH_CAMELLIA_128_CBC_SHA (rsa 2048)
- TLS_RSA_WITH_CAMELLIA_256_CBC_SHA (rsa 2048)
- TLS_RSA_WITH_IDEA_CBC_SHA (rsa 2048)
- TLS_RSA_WITH_RC4_128_MD5 (rsa 2048)
- TLS_RSA_WITH_RC4_128_SHA (rsa 2048)
- TLS_RSA_WITH_SEED_CBC_SHA (rsa 2048)

TLSv1.2 Ciphers:

- TLS_ECDHE_RSA_WITH_3DES_EDE_CBC_SHA (secp256r1)
- TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA (secp256r1)
- TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256 (secp256r1)
- TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256 (secp256r1)
- TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA (secp256r1)
- TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384 (secp256r1)
- TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384 (secp256r1)
- TLS_ECDHE_RSA_WITH_RC4_128_SHA (secp256r1)
- TLS_RSA_WITH_3DES_EDE_CBC_SHA (rsa 2048)
- TLS_RSA_WITH_AES_128_CBC_SHA (rsa 2048)
- TLS_RSA_WITH_AES_128_CBC_SHA256 (rsa 2048)
- TLS_RSA_WITH_AES_128_GCM_SHA256 (rsa 2048)
- TLS_RSA_WITH_AES_256_CBC_SHA (rsa 2048)
- TLS_RSA_WITH_AES_256_CBC_SHA256 (rsa 2048)
- TLS_RSA_WITH_AES_256_GCM_SHA384 (rsa 2048)
- TLS_RSA_WITH_CAMELLIA_128_CBC_SHA (rsa 2048)
- TLS_RSA_WITH_CAMELLIA_256_CBC_SHA (rsa 2048)
- TLS_RSA_WITH_IDEA_CBC_SHA (rsa 2048)
- TLS_RSA_WITH_RC4_128_MD5 (rsa 2048)
- TLS_RSA_WITH_RC4_128_SHA (rsa 2048)
- TLS_RSA_WITH_SEED_CBC_SHA (rsa 2048)

Configuring cipher suite selection using RACADM

To configure cipher suite selection using RACADM, use any one of the following commands:

- racadm set idraC.webServer.customCipherString ALL:!DHE-RSA-AES256-GCM-SHA384:!DHE-RSA-AES256-GCM-SHA384
- racadm set idraC.webServer.customCipherString ALL:-DHE-RSA-CAMELLIA256-SHA
- racadm set idraC.webServer.customCipherString ALL:!DHE-RSA-AES256-GCM-SHA384:!DHE-RSA-
- AES256-SHA256:+AES256-GCM-SHA384:-DHE-RSA-CAMELLIA256-SHA

For more information about these objects, see *iDRAC RACADM Command Line Interface Reference Guide* available at dell.com/ idracmanuals.

FIPS mode

FIPS is a computer security standard that United States government agencies and contractors must use. Starting from version iDRAC 2.40.40.40, iDRAC supports enabling FIPS mode.

iDRAC will be officially certified to support FIPS mode in the future.

Difference between FIPS-mode supported and FIPS-validated

Software that has been validated by completing the Cryptographic Module Validation Program is referred to as FIPS-validated. Because of the time it takes to complete FIPS-validation, not all versions of iDRAC are validated. For information about the latest status of FIPS-validation for iDRAC, see the Cryptographic Module Validation Program page on the NIST website.



Enabling FIPS Mode

CAUTION: Enabling FIPS mode resets iDRAC to factory-default settings. If you want to restore the settings, back up the server configuration profile (SCP) before you enable FIPS mode, and restore the SCP after iDRAC restarts.

(i) NOTE: If you reinstall or upgrade iDRAC firmware, FIPS mode gets disabled.

Enabling FIPS mode using web interface

- 1. On the iDRAC web interface, navigate to iDRAC Settings > Connectivity > Network > Network Settings > Advanced Network Settings.
- 2. In FIPS Mode, select Enabled and click Apply.

(i) NOTE: Enabling FIPS Mode resets iDRAC to the default settings.

- **3.** A message appears prompting you to confirm the change. Click **OK**. iDRAC restarts in FIPS mode. Wait for at least 60 seconds before you reconnect to iDRAC.
- 4. Install a trusted certificate for iDRAC.

(i) NOTE: The default SSL certificate is not allowed in FIPS mode.

NOTE: Some iDRAC interfaces, such as the standards-compliant implementations of IPMI and SNMP, do not support FIPS-compliance.

Enabling FIPS mode using RACADM

Use RACADM CLI to execute the following command:

```
racadm set iDRAC.Security.FIPSMode <Enable>
```

Disabling FIPS mode

To disable FIPS mode, you must reset iDRAC to the factory-default settings.

Configuring services

You can configure and enable the following services on iDRAC:

Local Configuration	Disable access to iDRAC configuration (from the host system) using Local RACADM and iDRAC Settings utility.
Web Server	Enable access to iDRAC web interface. If you disable the web interface, remote RACADM also gets disabled. Use local RACADM to re-enable the web server and remote RACADM.
SEKM Configuration	Enables secure enterprise key management functionality on iDRAC using a client server architecture.
SSH	Access iDRAC through firmware RACADM.
Remote RACADM	Remotely access iDRAC.
SNMP Agent	Enables support for SNMP queries (GET, GETNEXT, and GETBULK operations) in iDRAC.
Automated System Recovery Agent	Enable Last System Crash Screen.
Redfish	Enables support for Redfish RESTful API.

96 Configuring iDRAC



Enable VNC server with or without SSL encryption.



Configuring services using web interface

To configure the services using iDRAC Web interface:

- 1. In the iDRAC Web interface, go to **iDRAC Settings** > **Services**. The **Services** page is displayed.
- **2.** Specify the required information and click **Apply**.

For information about the various settings, see the *iDRAC Online Help*.

NOTE: Do not select the **Prevent this page from creating additional dialogs** check-box. Selecting this option prevents you from configuring services.

You can configure **SEKM** from iDRAC Settings page. Click **iDRAC Settings** > **Services** > **SEKM Configuration**.

(i) **NOTE:** For detailed step by step procedure for configuring SEKM, see the *iDRAC Online Help*.

NOTE: When **Security (Encryption)** mode is changed from **None** to **SEKM**, Real-Time job is not available. But it will be added to Staged job list. However, Real-Time job is successful when the mode is changed from **SEKM** to **None**.

Verify the following when changing the value of the **Username** Field in Client Certificate section on the KeySecure server (for ex: changing the value from **Common Name (CN)** to **User ID (UID)**)

- a. While using an existing account:
 - Verify in the iDRAC SSL certificate that, instead of the **Common Name** field, the **User name** field now matches the existing username on the KMS. If they don't, then you will have to set the username field and regenerate the SSL certificate again, get it signed on KMS and re-upload to iDRAC.
- **b.** While using a new user account:
 - Make sure the **User name** string matches the username field in the iDRAC SSL certificate.
 - If they don't match, then you will need to reconfigure the iDRAC KMS attributes Username and Password.
 - Once the certificate is verified to contain the username, then the only change that needs to be made is to change the key ownership from the old user to the new user to match the newly created KMS username.

While using Vormetric Data Security Manager as KMS, ensure that the Common Name (CN) field in iDRAC SSL certificate matches with the host name added to Vormetric Data Security Manager. Otherwise, the certificate may not import successfully.

() NOTE:

- Rekey option will be disabled when racadm sekm getstatus reports as Failed.
- SEKM only supports **Common name**, **User ID**, or **Organization Unit** for **User Name** field under Client certificate.
- If you are using a third party CA to sign the iDRAC CSR, ensure that the third party CA supports the value **UID** for **User Name** field in Client certificate. If it is not supported, use **Common Name** as the value for **User Name** field.
- If you are using Username and Password fields, ensure that KMS server supports those attributes.

NOTE: For KeySecure key management server,

- while creating an SSL certificate request, you must include the IP address of the key management server in **Subject** Alternative Name field
- the IP address must be in the following format: IP:xxx.xxx.xxx.xxx.

Configuring services using RACADM

To enable and configure services using RACADM, use the set command with the objects in the following object groups:

- iDRAC.LocalSecurity
- iDRAC.LocalSecurity
- iDRAC.SSH
- iDRAC.Webserver
- iDRAC.Racadm
- iDRAC.SNMP



For more information about these objects, see *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Enabling or disabling HTTPS redirection

If you do not want automatic redirection from HTTP to HTTPs due to certificate warning issue with default iDRAC certificate or as a temporary setting for debugging purpose, you can configure iDRAC such that redirection from http port (default is 80) to https port (default is 443) is disabled. By default, it is enabled. You have to log out and log in to iDRAC for this setting to take effect. When you disable this feature, a warning message is displayed.

You must have Configure iDRAC privilege to enable or disable HTTPS redirection.

An event is recorded in the Lifecycle Controller log file when this feature is enabled or disabled.

To disable the HTTP to HTTPS redirection:

racadm set iDRAC.Webserver.HttpsRedirection Disabled

To enable HTTP to HTTPS redirection:

racadm set iDRAC.Webserver.HttpsRedirection Enabled

To view the status of the HTTP to HTTPS redirection:

racadm get iDRAC.Webserver.HttpsRedirection

SEKM Functionalities

Following are the SEKM functionalities available in iDRAC:

- SEKM Key Purge Policy iDRAC provides a policy setting that allows you to configure iDRAC to purge old unused keys at the Key Management Server (KMS) when Rekey operation is performed. You can set iDRAC read-writable attribute KMSKeyPurgePolicy to one of the following values:
 - Keep All Keys This is the default setting and is the existing behavior where iDRAC leaves all the keys on the KMS untouched while performing Rekey operation.
 - Keep N and N-1 keys iDRAC deletes all keys at the KMS except the current (N) and previous key (N-1) when performing Rekey operation.
- 2. KMS Key Purge on SEKM Disable As part of the Secure Enterprise Key Manager (SEKM) solution, iDRAC allows you to disable SEKM on the iDRAC. Once SEKM is disabled, the keys generated by iDRAC at the KMS are unused and remain at the KMS. This feature is for allowing iDRAC to delete those keys when SEKM is disabled. iDRAC provides a new option "-purgeKMSKeys" to existing legacy command "racadm sekm disable" which will let you purge keys at the KMS when SEKM is disabled on iDRAC.

NOTE: If SEKM is already disabled and you want to purge old keys, you must re-enable SEKM, then disable passing in option -purgeKMSKeys.

- **3. Key Creation Policy** As part of this release, iDRAC has been pre-configured with a Key Creation Policy. Attribute KeyCreationPolicy is read only and set to "Key per iDRAC" value.
- iDRAC read-only attribute iDRAC.SEKM.KeyIdentifierN reports the Key Identifier created by the KMS.

racadm get iDRAC.SEKM.KeyIdentifierN

• iDRAC read-only attribute iDRAC.SEKM.KeyldentifierNMinusOne reports the previous Key Identifier after performing a Rekey operation.

racadm get iDRAC.SEKM.KeyIdentifierNMinusOne

- 4. SEKM Rekey iDRAC provides 2 options to rekey your SEKM solution, either Rekey iDRAC or PERC. It's recommended to rekey the iDRAC since this rekeys all SEKM Secure capable/Enabled devices.
 - SEKM iDRAC Rekey [Rekey on iDRAC.Embedded.1 FQDD] When performing racadm sekm rekey iDRAC.Embedded.1, all SEKM Secure capable/Enabled devices are Rekeyed with a new key from KMS and this is common key to all SEKM enabled devices. iDRAC Rekey operation can also be executed from iDRAC GUI- **iDRAC Settings** > **Services** > **SEKM Configuration** > **Rekey**. After executing this operation, the change in the Key can be validated by reading KeyldentifierN and KeyldentifierNMinusOne attributes.



 SEKM PERC Rekey (Rekey On Controller [Example RAID.Slot.1-1] FQDD) — When performing racadm sekm rekey <controller FQDD>, the corresponding SEKM enabled controller gets rekeyed to the currently active iDRAC common key created from KMS. Storage Controller Rekey operation can also be executed from iDRAC GUI- Storage > Controllers > <controller FQDD> > Actions > Edit > Security > Security(Encryption) > Rekey.

Using VNC client to manage remote server

You can use a standard open VNC client to manage the remote server using both desktop and mobile devices such as Dell Wyse PocketCloud. When servers in data centers stop functioning, the iDRAC or the operating system sends an alert to the console on the management station. The console sends an email or SMS to a mobile device with required information and launches VNC viewer application on the management station. This VNC viewer can connect to OS/Hypervisor on the server and provide access to keyboard, video and mouse of the host server to perform the necessary remediation. Before launching the VNC client, you must enable the VNC server and configure the VNC server settings in iDRAC such as password, VNC port number, SSL encryption, and the time out value. You can configure these settings using iDRAC Web interface or RACADM.

(i) NOTE: VNC feature is licensed and is available in the iDRAC Enterprise or Datacenter license.

You can choose from many VNC applications or Desktop clients such as the ones from RealVNC or Dell Wyse PocketCloud.

2 VNC client sessions can be activated at the same time. Second one is in Read-Only mode.

If a VNC session is active, you can only launch the Virtual Media using Launch Virtual Console and not the Virtual Console Viewer.

If video encryption is disabled, the VNC client starts RFB handshake directly, and a SSL handshake is not required. During VNC client handshake (RFB or SSL), if another VNC session is active or if a Virtual Console session is open, the new VNC client session is rejected. After completion of the initial handshake, VNC server disables Virtual Console and allows only Virtual Media. After termination of the VNC session, VNC server restores the original state of Virtual Console (enabled or disabled).

() NOTE:

- While launching a VNC session, if you get an RFB protocol error, change the VNC client settings to High quality and then relaunch the session.
- When iDRAC NIC is in shared mode and the host system is power cycled, the network connection is lost for a few seconds. During this time, if you perform any action in the active VNC client, the VNC session may close. You must wait for timeout (value configured for the VNC Server settings in the **Services** page in iDRAC Web interface) and then re-establish the VNC connection.
- If the VNC client window is minimized for more than 60 seconds, the client window closes. You must open a new VNC session. If you maximize the VNC client window within 60 seconds, you can continue to use it.

Configuring VNC server using iDRAC web interface

To configure the VNC server settings:

- 1. In the iDRAC Web interface, go to **Configuration** > **Virtual Console**. The **Virtual Console** page is displayed.
- 2. In the VNC Server section, enable the VNC server, specify the password, port number, and enable or disable SSL encryption.

For information about the fields, see the *iDRAC Online Help*.

3. Click **Apply**. The VNC server is configured.

Configuring VNC server using RACADM

To configure the VNC server, use the set command with the objects in ${\tt VNCserver}.$

For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.



Setting up VNC viewer with SSL encryption

While configuring the VNC server settings in iDRAC, if the **SSL Encryption** option was enabled, then the SSL tunnel application must be used along with the VNC Viewer to establish the SSL encrypted connection with iDRAC VNC server.

(i) NOTE: Most of the VNC clients do not have built-in SSL encryption support.

To configure the SSL tunnel application:

- 1. Configure SSL tunnel to accept connection on <localhost>:<localport number>. For example, 127.0.0.1:5930.
- Configure SSL tunnel to connect to <iDRAC IP address>:<VNC server port Number>. For example, 192.168.0.120:5901.
- 3. Start the tunnel application.

To establish connection with the iDRAC VNC server over the SSL encrypted channel, connect the VNC viewer to the localhost (link local IP address) and the local port number (127.0.0.1:<local port number>).

Setting up VNC viewer without SSL encryption

In general, all Remote Frame Buffer (RFB) compliant VNC Viewers connect to the VNC server using the iDRAC IP address and port number that is configured for the VNC server. If the SSL encryption option is disabled when configuring the VNC server settings in iDRAC, then to connect to the VNC Viewer do the following:

In the VNC Viewer dialog box, enter the iDRAC IP address and the VNC port number in the VNC Server field.

The format is <iDRAC IP address:VNC port number>

For example, if the iDRAC IP address is 192.168.0.120 and VNC port number is 5901, then enter 192.168.0.120:5901.

Configuring front panel display

You can configure the front panel LCD and LED display for the managed system.

- For rack and tower servers, two types of front panels are available:
- LCD front panel and System ID LED
- LED front panel and System ID LED

For blade servers, only the System ID LED is available on the server front panel since the blade chassis has the LCD.

Configuring LCD setting

You can set and display a default string such as iDRAC name, IP, and so on or a user-defined string on the LCD front panel of the managed system.

Configuring LCD setting using web interface

To configure the server LCD front panel display:

- 1. In iDRAC Web interface, go to Configurations > System Settings > Hardware Settings > Front Panel configuration.
- 2. In LCD Settings section, from the Set Home Message drop-down menu, select any of the following:
 - Service Tag (default)
 - Asset Tag
 - DRAC MAC Address
 - DRAC IPv4 Address
 - DRAC IPv6 Address
 - System Power
 - Ambient Temperature
 - System Model
 - Host Name
 - User Defined


None

If you select **User Defined**, enter the required message in the text box.

If you select **None**, home message is not displayed on the server LCD front panel.

- **3.** Enable Virtual Console indication (optional). If enabled, the Live Front Panel Feed section and the LCD panel on the server displays the Virtual console session active message when there is an active Virtual Console session.
- 4. Click Apply.

The server LCD front panel displays the configured home message.

Configuring LCD setting using RACADM

To configure the server LCD front panel display, use the objects in the System.LCD group.

For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Configuring LCD setting using iDRAC settings utility

To configure the server LCD front panel display:

- In the iDRAC Settings utility, go to Front Panel Security. The iDRAC Settings.Front Panel Security page is displayed.
- 2. Enable or disable the power button.
- **3.** Specify the following:
 - Access to the front panel
 - LCD message string
 - System power units, ambient temperature units, and error display
- 4. Enable or disable the virtual console indication.

For information about the options, see the *iDRAC Settings Utility Online Help*.

5. Click Back, click Finish, and then click Yes.

Configuring system ID LED setting

To identify a server, enable or disable System ID LED blinking on the managed system.

Configuring system ID LED setting using web interface

To configure the System ID LED display:

- 1. In iDRAC Web interface, go to Configuration > System Settings > Hardware Settings > Front Panel configuration. The System ID LED Settings page is displayed.
- 2. In System ID LED Settings section, select any of the following options to enable or disable LED blinking:
 - Blink Off
 - Blink On
 - Blink On 1 Day Timeout
 - Blink On 1 Week Timeout
 - Blink On 1 Month Timeout
- 3. Click Apply.

The LED blinking on the front panel is configured.

Configuring system ID LED setting using RACADM

To configure system ID LED, use the setled command.

For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Configuring iDRAC 101



Configuring time zone and NTP

You can configure the time zone on iDRAC and synchronize the iDRAC time using Network Time Protocol (NTP) instead of BIOS or host system times.

You must have Configure privilege to configure time zone or NTP settings.

Configuring time zone and NTP using iDRAC web interface

To configure time zone and NTP using iDRAC web interface:

- Go to iDRAC Settings > Settings > Time zone and NTP Settings. The Time zone and NTP page is displayed.
- To configure the time zone, from the Time Zone drop-down menu, select the required time zone, and then click Apply.
- 3. To configure NTP, enable NTP, enter the NTP server addresses, and then click **Apply**.

For information about the fields, see iDRAC Online Help.

Configuring time zone and NTP using RACADM

To configure time zone and NTP, use the set command with the objects in the iDRAC.Time and iDRAC.NTPConfigGroup group.

For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

() NOTE: iDRAC syncs the time with the host (local time). Hence it is recommended to configure both iDRAC and host with the same time zone so that the time sync is proper. If you want to change a time zone, you need to change it on both host and iDRAC and then the host needs to reboot.

Setting first boot device

You can set the first boot device for the next boot only or for all subsequent reboots. If you set the device to be used for all subsequent boots, it remains as the first boot device in the BIOS boot order until it is changed again either from the iDRAC web interface or from the BIOS boot sequence.

You can set the first boot device to one of the following:

- Normal Boot
- PXE
- BIOS Setup
- Local Floppy/Primary Removable Media
- Local CD/DVD
- Hard Drive
- Virtual Floppy
- Virtual CD/DVD/ISO
- Local SD Card
- Lifecycle Controller
- BIOS Boot Manager
- UEFI Device Path
- UEFI HTTP

() NOTE:

- BIOS Setup (F2), Lifecycle Controller (F10), and BIOS Boot Manager (F11) cannot be set as permanent boot device.
- The first boot device setting in iDRAC Web Interface overrides the System BIOS boot settings.

Setting first boot device using web interface

To set the first boot device using iDRAC Web interface:

102 Configuring iDRAC



- Go to Configuration > System Settings > Hardware Settings > First Boot Device. The First Boot Device page is displayed.
- Select the required first boot device from the drop-down list, and click Apply. The system boots from the selected device for subsequent reboots.
- To boot from the selected device only once on the next boot, select Boot Once. Thereafter, the system boots from the first boot device in the BIOS boot order.

For more information about the options, see the *iDRAC Online Help*.

Setting first boot device using RACADM

- To set the first boot device, use the iDRAC.ServerBoot.FirstBootDevice object.
- To enable boot once for a device, use the iDRAC.ServerBoot.BootOnce object.

For more information about these objects, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Setting first boot device using virtual console

You can select the device to boot from as the server is being viewed in the Virtual Console viewer before the server runs through its boot-up sequence. Boot-once is supported by all devices listed in Setting first boot device.

To set the first boot device using Virtual Console:

- 1. Launch Virtual Console.
- 2. In the Virtual Console Viewer, from the Next Boot menu, set the required device as the first boot device.

Enabling last crash screen

To troubleshoot the cause of a crash on the managed system, you can capture the system crash image using iDRAC. **NOTE:** For information about Server Administrator, see the *OpenManage Installation Guide* available at https://www.dell.com/openmanagemanuals.

The host system should have Windows Operating system to use this feature.

() NOTE:

- This feature is not applicable on Linux system.
- This feature is independent of any agents or attributes.

Enabling or disabling OS to iDRAC Pass-through

In servers that have Network Daughter Card (NDC) or embedded LAN On Motherboard (LOM) devices, you can enable the OS to iDRAC Pass-through feature. This feature provides a high-speed bi-directional in-band communication between iDRAC and the host operating system through a shared LOM, a dedicated NIC, or through the USB NIC. This feature is available for iDRAC Enterprise or Datacenter license.

NOTE: iDRAC Service Module (iSM) provides more features for managing iDRAC through the operating system. For more information, see the iDRAC Service Module User's Guide available at www.dell.com/idracservicemodule.

When enabled through dedicated NIC, you can launch the browser in the host operating system and then access the iDRAC Web interface. The dedicated NIC for the blade servers is through the Chassis Management Controller.

Switching between dedicated NIC or shared LOM does not require a reboot or reset of the host operating system or iDRAC.

You can enable this channel using:

- iDRAC web interface
- RACADM or WSMan (post operating system environment)
- iDRAC Settings utility (pre-operating system environment)

If the network configuration is changed through iDRAC Web interface, you must wait for at least 10 seconds before enabling OS to iDRAC Pass-through.

Configuring iDRAC 103



If you are configuring the server using a Server Configuration Profile through RACADM, WSMan or Redfish and if the network settings are changed in this file, then you must wait for 15 seconds to either enable OS to iDRAC Pass-through feature or set the OS Host IP address.

Before enabling OS to iDRAC Pass-through, make sure that:

- iDRAC is configured to use dedicated NIC or shared mode (that is, NIC selection is assigned to one of the LOMs).
- Host operating system and iDRAC are in the same subnet and same VLAN.
- Host operating system IP address is configured.
- A card that supports OS to iDRAC Pass-through capability is installed.
- You have the Configure privilege.

When you enable this feature:

- In shared mode, the host operating system's IP address is used.
- In dedicated mode, you must provide a valid IP address of the host operating system. If more than one LOM is active, enter the first LOM's IP address.

If the OS to iDRAC Pass-through feature does not work after it is enabled, ensure that you check the following:

- The iDRAC dedicated NIC cable is connected properly.
- At least one LOM is active.

NOTE: Use the default IP address. Ensure that the IP address of the USB NIC interface is not in the same network subnet as the iDRAC or host OS IP addresses. If this IP address conflicts with an IP address of other interfaces of the host system or the local network, you must change it.

- **NOTE:** If you launch iDRAC Service Module while USB NIC is in disabled state, the iDRAC Service Module changes the USB NIC IP address to 169.254.0.1.
- **NOTE:** Do not use 169.254.0.3 and 169.254.0.4 IP addresses. These IP addresses are reserved for the USB NIC port on the front panel when an A/A cable is used.
- () NOTE: iDRAC may not be accessible from the host server using LOM-Passthrough when NIC teaming is enabled. Then, iDRAC can be accessed from the host server OS using the iDRAC USB NIC or through the external network, via the iDRAC dedicated NIC.

Supported cards for OS to iDRAC Pass-through

The following table provides a list of cards that support the OS to iDRAC Pass-through feature using LOM.

Table 15. OS to iDRAC Pass-through using LOM — supported cards

Category	Manufacturer	Туре
NDC	Broadcom	• 5720 QP rNDC 1G BASE-T
	Intel	• x520/i350 QP rNDC 1G BASE-T

In-built LOM cards also support the OS to iDRAC pass-through feature.

Supported operating systems for USB NIC

The operating systems supported for USB NIC are:

- Server 2012 R2 Foundation Edition
- Server 2012 R2 Essentials Edition
- Server 2012 R2 Standard Edition
- Server 2012 R2 Datacenter Edition
- Server 2012 for Embedded Systems (Base and R2 w/ SP1)
- Server 2016 Essentials Edition
- Server 2016 Standard Edition
- Server 2016 Datacenter Edition
- RHEL 7.3
- RHEL 6.9



- SLES 12 SP2
- ESXi 6.0 U3
- vSphere 2016
- XenServer 7.1

For Linux operating systems, configure the USB NIC as DHCP on the host operating system before enabling USB NIC.

For vSphere, you must install the VIB file before enabling USB NIC.

NOTE: To configure USB NIC as DHCP in Linux operating system or XenServer, refer to the operating system or hypervisor documentation.

Installing VIB file

For vSphere operating systems, before enabling the USB NIC, you must install the VIB file.

To install the VIB file:

- 1. Using Win-SCP, copy the VIB file to /tmp/ folder of the ESX-i host operating system.
- 2. Go to the ESXi prompt and run the following command:

```
esxcli software vib install -v /tmp/ iDRAC USB NIC-1.0.0-799733X03.vib --no-sig-check
```

The output is:

```
Message: The update completed successfully, but the system needs to be rebooted for
the changes to be effective.
Reboot Required: true
VIBs Installed: Dell_bootbank_iDRAC_USB_NIC_1.0.0-799733X03
VIBs Removed:
VIBs Skipped:
```

- 3. Reboot the server.
- At the ESXi prompt, run the command: esxcfg-vmknic -1. The output displays the usb0 entry.

Enabling or disabling OS to iDRAC Pass-through using web interface

To enable OS to iDRAC Pass-through using Web interface:

- Go to iDRAC Settings > Connectivity > Network > OS to iDRAC Pass-through. The OS to iDRAC Pass-through page is displayed.
- 2. Change the State to Enabled.
- **3.** Select any of the following options for Pass-through Mode:
 - LOM The OS to iDRAC pass-through link between the iDRAC and the host operating system is established through the LOM or NDC.
 - **USB NIC** The OS to iDRAC pass-through link between the iDRAC and the host operating system is established through the internal USB bus.
 - (i) NOTE: If you set the pass-through mode to LOM, ensure that:
 - OS and iDRAC are on the same subnet
 - NIC selection in Network Settings is set to LOM
- 4. If the server is connected in shared LOM mode, then the OS IP Address field is disabled.

NOTE: If the VLAN is enabled on the iDRAC, the LOM-Passthrough will only function in shared LOM mode with VLAN tagging configured on the host.

() NOTE:

• When Pass-through mode is set to LOM, it is not possible to launch iDRAC from host OS after cold boot.



- We have purposefully removed the LOM Pass-through using Dedicated mode feature.
- 5. If you select **USB NIC** as the pass-through configuration, enter the IP address of the USB NIC.

The default value is 169.254.1.1. It is recommended to use the default IP address. However, if this IP address conflicts with an IP address of other interfaces of the host system or the local network, you must change it.

Do not enter 169.254.0.3 and 169.254.0.4 IPs. These IPs are reserved for the USB NIC port on the front panel when an A/A cable is used.

NOTE: If IPv6 is preferred, the default address is fde1:53ba:e9a0:de11::1. If needed, this address can be modified in the idrac.OS-BMC.UsbNicULA setting. If IPv6 is not wanted on the USB-NIC, it can be disabled by changing the address to "::"

- 6. Click Apply.
- 7. Click **Test Network Configuration** to check if the IP is accessible and the link is established between the iDRAC and the host operating system.

Enabling or disabling OS to iDRAC Pass-through using RACADM

To enable or disable OS to iDRAC Pass-through using RACADM, use the objects in the iDRAC.OS-BMC group. For more information, see theiDRAC Attribute Registry available at https://www.dell.com/idracmanuals.

Enabling or disabling OS to iDRAC Pass-through using iDRAC settings utility

To enable or disable OS to iDRAC Pass-through using iDRAC Settings Utility:

- In the iDRAC Settings utility, go to Communications Permissions. The iDRAC Settings.Communications Permissions page is displayed.
- 2. Select any of the following options to enable OS to iDRAC pass-through:
 - LOM The OS to iDRAC pass-through link between the iDRAC and the host operating system is established through the LOM or NDC.
 - **USB NIC** The OS to iDRAC pass-through link between the iDRAC and the host operating system is established through the internal USB bus.
 - () NOTE: If you set the pass-through mode to LOM, ensure that:
 - OS and iDRAC are on the same subnet
 - NIC selection in Network Settings is set to a LOM
 - To disable this feature, select **Disabled**.

NOTE: The LOM option can be selected only of the card supports OS to iDRAC pass-through capability. Else, this option is grayed-out.

3. If you select **LOM** as the pass-through configuration, and if the server is connected using dedicated mode, enter the IPv4 address of the operating system.

i) NOTE: If the server is connected in shared LOM mode, then the OS IP Address field is disabled.

4. If you select USB NIC as the pass-through configuration, enter the IP address of the USB NIC.

The default value is 169.254.1.1. However, if this IP address conflicts with an IP address of other interfaces of the host system or the local network, you must change it. Do not enter 169.254.0.3 and 169.254.0.4 IPs. These IPs are reserved for the USB NIC port on the front panel when an A/A cable is used.

() NOTE: If IPv6 is preferred, the default address is fde1:53ba:e9a0:de11::1. If needed, this address can be modified in the idrac.OS-BMC.UsbNicULA setting. If IPv6 is not wanted on the USB-NIC, it can be disabled by changing the address to "::"

5. Click **Back**, click **Finish**, and then click **Yes**.

The details are saved.



Obtaining certificates

The following table lists the types of certificates based on the login type.

Table 16. Types of certificate based on login type

Login Type	Certificate Type	How to Obtain
Single Sign-on using Active Directory	Trusted CA certificate	Generate a CSR and get it signed from a Certificate Authority
		SHA-2 certificates are also supported.
Smart Card login as a local or Active Directory user	 User certificate Trusted CA certificate 	 User Certificate — Export the smart card user certificate as Base64-encoded file using the card management software provided by the smart card vendor. Trusted CA certificate — This certificate is issued by a CA. SHA-2 certificates are also supported.
Active Directory user login	Trusted CA certificate	This certificate is issued by a CA. SHA-2 certificates are also
		supported.
Local User login	SSL Certificate	Generate a CSR and get it signed from a trusted CA
		(i) NOTE: iDRAC ships with a default self-signed SSL server certificate. The iDRAC Web server, Virtual Media, and Virtual Console use this certificate.
		SHA-2 certificates are also supported.

SSL server certificates

iDRAC includes a web server that is configured to use the industry-standard SSL security protocol to transfer encrypted data over a network. An SSL encryption option is provided to disable weak ciphers. Built upon asymmetric encryption technology, SSL is widely accepted for providing authenticated and encrypted communication between clients and servers to prevent eavesdropping across a network.

An SSL-enabled system can perform the following tasks:

- Authenticate itself to an SSL-enabled client
- Allow the two systems to establish an encrypted connection
- **NOTE:** If SSL encryption is set to 256-bit or higher and 168–bit or higher, the cryptography settings for your virtual machine environment (JVM, IcedTea) may require installing the Unlimited Strength Java Cryptography Extension Policy Files to permit usage of iDRAC plugins such as vConsole with this level of encryption. For information about installing the policy files, see the documentation for Java.

iDRAC Web server has a Dell self-signed unique SSL digital certificate by default. You can replace the default SSL certificate with a certificate signed by a well-known Certificate Authority (CA). A Certificate Authority is a business entity that is recognized in the Information Technology industry for meeting high standards of reliable screening, identification, and other important security criteria. Examples of CAs include Thawte and VeriSign. To initiate the process of obtaining a CA-signed certificate, use either iDRAC Web interface or RACADM interface to generate a Certificate Signing Request (CSR) with your



company's information. Then, submit the generated CSR to a CA such as VeriSign or Thawte. The CA can be a root CA or an intermediate CA. After you receive the CA-signed SSL certificate, upload this to iDRAC.

For each iDRAC to be trusted by the management station, that iDRAC's SSL certificate must be placed in the management station's certificate store. Once the SSL certificate is installed on the management stations, supported browsers can access iDRAC without certificate warnings.

You can also upload a custom signing certificate to sign the SSL certificate, rather than relying on the default signing certificate for this function. By importing one custom signing certificate into all management stations, all the iDRACs using the custom signing certificate are trusted. If a custom signing certificate is uploaded when a custom SSL certificate is already in-use, then the custom SSL certificate is disabled and a one-time auto-generated SSL certificate, signed with the custom signing certificate, is used. You can download the custom signing certificate, iDRAC resets and auto-generates a new self-signed SSL certificate. If a self-signed certificate is regenerated, then the trust must be re-established between that iDRAC and the management workstation. Auto-generated SSL certificates are self-signed and have an expiration date of seven years and one day and a start date of one day in the past (for different time zone settings on management stations and the iDRAC).

The iDRAC Web server SSL certificate supports the asterisk character (*) as part of the left-most component of the Common Name when generating a Certificate Signing Request (CSR). For example, *.qa.com, or *.company.qa.com. This is called a wildcard certificate. If a wildcard CSR is generated outside of iDRAC, you can have a signed single wildcard SSL certificate that you can upload for multiple iDRACs and all the iDRACs are trusted by the supported browsers. While connecting to iDRAC Web interface using a supported browser that supports a wildcard certificate, the iDRAC is trusted by the browser. While launching viewers, the iDRACs are trusted by the viewer clients.

Generating a new certificate signing request

A CSR is a digital request to a Certificate Authority (CA) for a SSL server certificate. SSL server certificates allow clients of the server to trust the identity of the server and to negotiate an encrypted session with the server.

After the CA receives a CSR, they review and verify the information the CSR contains. If the applicant meets the CA's security standards, the CA issues a digitally-signed SSL server certificate that uniquely identifies the applicant's server when it establishes SSL connections with browsers running on management stations.

After the CA approves the CSR and issues the SSL server certificate, it can be uploaded to iDRAC. The information used to generate the CSR, stored on the iDRAC firmware, must match the information contained in the SSL server certificate, that is, the certificate must have been generated using the CSR created by iDRAC.

Generating CSR using web interface

To generate a new CSR:

NOTE: Each new CSR overwrites any previous CSR data stored in the firmware. The information in the CSR must match the information in the SSL server certificate. Else, iDRAC does not accept the certificate.

- In the iDRAC Web interface, go to iDRAC Settings > Services > Web Server > SSL certificate, select Generate Certificate Signing Request (CSR) and click Next. The Generate a New Certificate Signing Request page is displayed.
- **2.** Enter a value for each CSR attribute. For more information, see *iDRAC Online Help*.
- 3. Click Generate.

A new CSR is generated. Save it to the management station.

Generating CSR using RACADM

To generate a CSR using RACADM, use the set command with the objects in the iDRAC.Security group, and then use the sslcsrgen command to generate the CSR.

For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.



Automatic Certificate Enrollment

In iDRAC, Automatic certificate enrollment feature enables you for automatic installation and renewal of certificates used by the web server. When this feature is enabled, the existing web server certificate is replaced by a new certificate.

() NOTE:

- Automatic certificate enrollment is a licensed feature and requires Datacenter license.
- Valid NDES (Network Device Enrollment Service) setup is required for issuing the server certificate.

Following are the automatic certificate enrollment configuration parameters:

- Enable / Disable
- SCEP server URL
- Challenge password

(i) **NOTE:** For more information on these parameters, see *iDRAC Online Help*.

Following are the available status for Automatic certificate enrollment:

- Enrolled Automatic certificate enrollment is enabled. Certificate is monitored and new certificate can be issued on expiry.
- Enrolling Intermediate state after Automatic certificate enrollment is enabled.
- Error Problem encountered with NDES server.
- None Default.

NOTE: When you enable Automatic certificate enrollment, web server is restarted and all existing web sessions are logged out.

Uploading server certificate

After generating a CSR, you can upload the signed SSL server certificate to the iDRAC firmware. iDRAC must be reset to apply the certificate. iDRAC accepts only X509, Base 64 encoded Web server certificates. SHA-2 certificates are also supported.

CAUTION: During reset, iDRAC is not available for a few minutes.

Uploading server certificate using web interface

To upload the SSL server certificate:

- In the iDRAC Web interface, go to iDRAC Settings > Connectivity > SSL > SSL certificate, select Upload Server Certificate and click Next.
 - The **Certificate Upload** page is displayed.
- 2. Under File Path, click Browse and select the certificate on the management station.
- 3. Click Apply.
 - The SSL server certificate is uploaded to iDRAC.
- 4. A pop-up message is displayed asking you to reset iDRAC immediately or at a later time. Click **Reset iDRAC** or **Reset iDRAC** immediately or at a later time. Click **Reset iDRAC** or **Reset iDRAC** is a second sec

iDRAC resets and the new certificate is applied. The iDRAC is not available for a few minutes during the reset.

(i) NOTE: You must reset iDRAC to apply the new certificate. Until iDRAC is reset, the existing certificate is active.

Uploading server certificate using RACADM

To upload the SSL server certificate, use the sslcertupload command. For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

If the CSR is generated outside of iDRAC with a private key available, then to upload the certificate to iDRAC:

- 1. Send the CSR to a well-known root CA. CA signs the CSR and the CSR becomes a valid certificate.
- 2. Upload the private key using the remote racadm sslkeyupload command.
- **3.** Upload the signed certificate to iDRAC using the remote racadm sslcertupload command. The new certificate is uploaded iDRAC. A message is displayed asking you to reset iDRAC.



4. Run the racadm racreset command to reset iDRAC.

iDRAC resets and the new certificate is applied. The iDRAC is not available for a few minutes during the reset.

(i) NOTE: You must reset iDRAC to apply the new certificate. Until iDRAC is reset, the existing certificate is active.

Viewing server certificate

You can view the SSL server certificate that is currently being used in iDRAC.

Viewing server certificate using web interface

In the iDRAC Web interface, go to **iDRAC Settings** > **Services** > **Web Server** > **SSL certificate**. The **SSL** page displays the SSL server certificate that is currently in use at the top of the page.

Viewing server certificate using RACADM

To view the SSL server certificate, use the sslcertview command.

For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Uploading custom signing certificate

You can upload a custom signing certificate to sign the SSL certificate. SHA-2 certificates are also supported.

Uploading custom signing certificate using web interface

To upload the custom signing certificate using iDRAC web interface:

- Go to iDRAC Settings > Connectivity > SSL. The SSL page is displayed.
- 2. Under Custom SSL Certificate Signing Certificate, click Upload Signing Certificate. The Upload Custom SSL Certificate Signing Certificate page is displayed.
- **3.** Click **Choose File** and select the custom SSL certificate signing certificate file. Only Public-Key Cryptography Standards #12 (PKCS #12) compliant certificate is supported.
- 4. If the certificate is password protected, in the PKCS#12 Password field, enter the password.
- 5. Click Apply.
- The certificate is uploaded to iDRAC.
- 6. A pop-up message is displayed asking you to reset iDRAC immediately or at a later time. Click **Reset iDRAC** or **Reset iDRAC** immediately or at a later time. Click **Reset iDRAC** or **Reset iDRAC** immediately or at a later time. Click **Reset iDRAC** or **Reset iDRAC** immediately or at a later time. Click **Reset iDRAC** or **Reset iDRAC** immediately or at a later time. Click **Reset iDRAC** or **Reset iDRAC** immediately or at a later time. Click **Reset iDRAC** or **Reset iDRAC** immediately or at a later time. Click **Reset iDRAC** or **Reset iDRAC** immediately or at a later time. Click **Reset iDRAC** or **Reset iDRAC** immediately or at a later time. Click **Reset iDRAC** or **Reset iDRAC** immediately or at a later time. Click **Reset iDRAC** or **Reset iDRAC** immediately or at a later time. Click **Reset iDRAC** or **Reset iDRAC** immediately or at a later time. Click **Reset iDRAC** immediately or at a later time.

After iDRAC resets, the new certificate is applied. The iDRAC is not available for a few minutes during the reset.

(i) NOTE: You must reset iDRAC to apply the new certificate. Until iDRAC is reset, the existing certificate is active.

Uploading custom SSL certificate signing certificate using RACADM

To upload the custom SSL certificate signing certificate using RACADM, use the sslcertupload command, and then use the racreset command to reset iDRAC.

For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Downloading custom SSL certificate signing certificate

You can download the custom signing certificate using iDRAC Web interface or RACADM.



Downloading custom signing certificate

To download the custom signing certificate using iDRAC Web interface:

- 1. Go to **iDRAC Settings** > **Connectivity** > **SSL**. The **SSL** page is displayed.
- 2. Under Custom SSL Certificate Signing Certificate, select Download Custom SSL Certificate Signing Certificate and click Next.

A pop-up message is displayed that allows you to save the custom signing certificate to a location of your choice.

Downloading custom SSL certificate signing certificate using RACADM

To download the custom SSL certificate signing certificate, use the sslcertdownload subcommand. For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Deleting custom SSL certificate signing certificate

You can also delete an existing custom signing certificate using iDRAC Web interface or RACADM.

Deleting custom signing certificate using iDRAC web interface

To delete the custom signing certificate using iDRAC web interface:

- Go to iDRAC Settings > Connectivity > SSL. The SSL page is displayed.
- 2. Under Custom SSL Certificate Signing Certificate, select Delete Custom SSL Certificate Signing Certificate and click Next.
- 3. A pop-up message is displayed asking you to reset iDRAC immediately or at a later time. Click **Reset iDRAC** or **Reset iDRAC Later** as required.

After iDRAC resets, a new self-signed certificate is generated.

Deleting custom SSL certificate signing certificate using RACADM

To delete the custom SSL certificate signing certificate using RACADM, use the sslcertdelete subcommand. Then, use the racreset command to reset iDRAC.

For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Configuring multiple iDRACs using RACADM

You can configure one or more iDRACs with identical properties using RACADM. When you query a specific iDRAC using its group ID and object ID, RACADM creates a configuration file from the retrieved information. Import the file to other iDRACs to identically configure them.

() NOTE:

- The configuration file contains information that is applicable for the particular server. The information is organized under various object groups.
- Some configuration files contain unique iDRAC information, such as the static IP address, that you must modify before you import the file into other iDRACs.

You can also use the System Configuration Profile (SCP) to configure multiple iDRACs using RACADM. SCP file contains the component configuration information. You can use this file to apply the configuration for BIOS, iDRAC, RAID, and NIC by importing the file into a target system. For more information, see *XML Configuration Workflow* white paper available at https://www.dell.com/manuals.

To configure multiple iDRACs using the configuration file:



1. Query the target iDRAC that contains the required configuration using the following command:.

racadm get -f <file name>.xml -t xml -c iDRAC.Embedded.1

The command requests the iDRAC configuration and generates the configuration file.

NOTE: Redirecting the iDRAC configuration to a file using get -f is only supported with the local and remote RACADM interfaces.

(i) NOTE: The generated configuration file does not contain user passwords.

The get command displays all configuration properties in a group (specified by group name and index) and all configuration properties for a user.

2. Modify the configuration file using a text editor, if required.

NOTE: It is recommended that you edit this file with a simple text editor. The RACADM utility uses an ASCII text parser. Any formatting confuses the parser, which may corrupt the RACADM database.

3. On the target iDRAC, use the following command to modify the settings:

```
racadm set -f <file name>.xml -t xml
```

This loads the information into the other iDRAC. You can use set command to synchronize the user and password database with Server Administrator.

4. Reset the target iDRAC using the command: racadm racreset

Disabling access to modify iDRAC configuration settings on host system

You can disable access to modify the iDRAC configuration settings through Local RACADM or iDRAC Settings utility. However, you can view these configuration settings. To do this:

- 1. In iDRAC Web interface, go to iDRAC Settings > Services > Local Configurations.
- 2. Select one or both of the following:
 - **Disable the iDRAC Local Configuration using iDRAC Settings** Disables access to modify the configuration settings in iDRAC Settings utility.
 - **Disable the iDRAC Local Configuration using RACADM** Disables access to modify the configuration settings in Local RACADM.
- 3. Click Apply.

NOTE: If access is disabled, you cannot use Server Administrator or IPMITool to perform iDRAC configurations. However, you can use IPMI Over LAN.



Delegated Authorization using OAuth 2.0

The Delegated Authorization feature allows a user or console to access iDRAC API using OAuth 2.0 JSON Web Tokens (JWT) that the user or console first obtains from an Authorization Server. Once an OAuth JWT has been retrieved, the user or console may use it to invoke iDRAC API. This circumvents the need for specifying username and password to access the API.

NOTE: This feature is only available for DataCenter license. You need to have Configure iDRAC or Configure Users privilege to use this feature.

iDRAC supports configuration of up to 2 Authorization Servers. The configuration requires a user to specify the following Authorization Server details:

- Name A string to identify the Authorization Server on the iDRAC.
- Metadata URL The OpenID Connect compliant URL as advertised by the server.
- HTTPS certificate The server public key the iDRAC should use to communicate with the server.
- Offline Key The JWK set document for the Authorization Server.
- Offline Issuer The issuer string as used in tokens issued by the Authorization Server.

For Online configuration:

- When configuring an Authorization Server, the iDRAC administrator needs to ensure that the iDRAC has online network access to the Authorization Server.
- If iDRAC cannot access the Authorization Server, the configuration fails and a subsequent attempt to access the iDRAC API fails even though a valid token is presented.

For offline configuration:

• iDRAC does not need to communicate with the Auth server, but instead it is configures with the metadata details that it has downloaded offline. When configured offline, iDRAC has public portion of the signing keys and can validate the token without a network connection to the Auth server.



Viewing iDRAC and managed system information

You can view iDRAC and managed system health and properties, hardware and firmware inventory, sensor health, storage devices, network devices, and view and terminate user sessions. For blade servers, you can also view the Flex Address or Remote-Assigned Address (applicable only for MX platforms).

Topics:

- Viewing managed system health and properties
- Configuring Asset Tracking
- Viewing system inventory
- Viewing sensor information
- Monitoring performance index of CPU, memory, and input output modules
- Idle Server Detection
- GPU (Accelerators) Management
- Checking the system for Fresh Air compliance
- Viewing historical temperature data
- Viewing network interfaces available on host OS
- Viewing network interfaces available on host OS using RACADM
- Viewing FlexAddress mezzanine card fabric connections
- Viewing or terminating iDRAC sessions

Viewing managed system health and properties

When you log in to the iDRAC web interface, the **System Summary** page allows you to view the managed system's health, basic iDRAC information, preview the virtual console, add and view work notes, and quickly launch tasks such as power on or off, power cycle, view logs, update and rollback firmware, switch on or switch off the front panel LED, and reset iDRAC.

To access the **System Summary** page, go to **System** > **Overview** > **Summary**. The **System Summary** page is displayed. For more information, see the *iDRAC Online Help*.

You can also view the basic system summary information using the iDRAC Settings utility. To do this, in iDRAC Settings utility, go to **System Summary**. The **iDRAC Settings System Summary** page is displayed. For more information, see the *iDRAC Settings Utility Online Help*.

Configuring Asset Tracking

The Asset Tracking feature in iDRAC provides you the ability to configure various attributes that are related to your server. This includes information such as acquisition, warranty, service, and so on.

NOTE: Asset Tracking in iDRAC is similar to the Asset Tag feature in OpenManage Server Administrator. However, the attribute information has to be entered separately in both these tools for them to report the relevant Asset data.

To configure Asset Tracking:

- 1. In the iDRAC interface, go to **Configuration** > **Asset Tracking**.
- 2. Click Add Custom Assets to add any additional attributes which are not specified by default on this page.
- 3. Enter all the relevant information of your server asset and click Apply.
- 4. To view the Asset Tracking Report, go to System > Details > Asset Tracking.



Viewing system inventory

You can view information about the hardware and firmware components installed on the managed system. To do this, in iDRAC web interface, go to **System** > **Inventories**. For information about the displayed properties, see the *iDRAC Online Help*.

The Hardware Inventory section displays the information for the following components available on the managed system:

- iDRAC
- RAID controller
- Batteries
- CPUs
- DIMMs
- HDDs
- Backplanes
- Network Interface Cards (integrated and embedded)
- Video card
- SD card
- Power Supply Units (PSUs)
- Fans
- Fibre Channel HBAs
- USB
- NVMe PCIe SSD devices

The Firmware Inventory section displays the firmware version for the following components:

- BIOS
- Lifecycle Controller
- iDRAC
- OS driver pack
- 32-bit diagnostics
- System CPLD
- PERC controllers
- Batteries
- Physical disks
- Power supply
- NIC
- Fibre Channel
- Backplane
- Enclosure
- PCle SSDs

(i) NOTE:

- Software inventory displays only the last 4 bytes of the firmware version and the Release date information. For example, if the firmware version is FLVDL06, the firmware inventory displays DL06.
- When collecting software inventory using Redfish interface, the Release date information is displayed only for components which support rollback.

NOTE: On the Dell PowerEdge FX2/FX2s servers, the naming convention of the CMC version displayed in the iDRAC GUI differs from that on the CMC GUI. However, the version remains the same.

When you replace any hardware component or update the firmware versions, make sure to enable and run the **Collect System Inventory on Reboot** (CSIOR) option to collect the system inventory on reboot. After a few minutes, log in to iDRAC, and navigate to the **System Inventory** page to view the details. It may take up to 5 minutes for the information to be available depending on the hardware installed on the server.

(i) NOTE: CSIOR option is enabled by default.

NOTE: Configuration changes and firmware updates that are made within the operating system may not reflect properly in the inventory until you perform a server restart.

Click **Export** to export the hardware inventory in an XML format and save it to a location of your choice.



Viewing sensor information

The following sensors help to monitor the health of the managed system:

• **Batteries** — Provides information about the batteries on the system board CMOS and storage RAID On Motherboard (ROMB).

(i) NOTE: The Storage ROMB battery settings are available only if the system has a ROMB with a battery.

- Fan (available only for rack and tower servers) Provides information about the system fans fan redundancy and fans list that display fan speed and threshold values.
- **CPU** Indicates the health and state of the CPUs in the managed system. It also reports processor automatic throttling and predictive failure.
- **Memory** Indicates the health and state of the Dual In-line Memory Modules (DIMMs) present in the managed system.
- Intrusion Provides information about the chassis.
- **Power Supplies** (available only for rack and tower servers) Provides information about the power supplies and the power supply redundancy status.

(i) NOTE: If there is only one power supply in the system, the power supply redundancy is set to **Disabled**.

- **Removable Flash Media** Provides information about the Internal SD Modules; vFlash and Internal Dual SD Module (IDSDM).
 - When IDSDM redundancy is enabled, the following IDSDM sensor status is displayed IDSDM Redundancy Status, IDSDM SD1, IDSDM SD2. When redundancy is disabled, only IDSDM SD1 is displayed.
 - If IDSDM redundancy is initially disabled when the system is powered on or after an iDRAC reset, the IDSDM SD1 sensor status is displayed only after a card is inserted.
 - If IDSDM redundancy is enabled with two SD cards present in the IDSDM, and the status of one SD card is online while the status of the other card is offline. A system reboot is required to restore redundancy between the two SD cards in the IDSDM. After the redundancy is restored, the status of both the SD cards in the IDSDM is online.
 - During the rebuilding operation to restore redundancy between two SD cards present in the IDSDM, the IDSDM status is not displayed since the IDSDM sensors are powered off.

(i) **NOTE:** If the host system is rebooted during IDSDM rebuild operation, the iDRAC does not display the IDSDM information. To resolve this, rebuild IDSDM again or reset the iDRAC.

- System Event Logs (SEL) for a write-protected or corrupt SD card in the IDSDM module are not repeated until they are cleared by replacing the SD card with a writable or good SD card, respectively.
- **NOTE:** When iDRAC firmware is updated from versions prior to 3.30.30.30, the iDRAC need to be reset to defaults for IDSDM settings to appear in the Server Administrator's Platform Event Filter.
- **Temperature** Provides information about the system board inlet temperature and exhaust temperature (only applies to rack servers). The temperature probe indicates whether the status of the probe is within the preset warning and critical threshold value.
- Voltage Indicates the status and reading of the voltage sensors on various system components.

The following table provides information about viewing the sensor information using iDRAC web interface and RACADM. For information about the properties that are displayed on the web interface, see the *iDRAC Online Help*.

(i) NOTE: The Hardware Overview page displays data only for sensors present on your system.

Table 17. Sensor information using web interface and RACADM

View sensor information For	Using web interface	Using RACADM
Batteries	Dashboard > System Health > Batteries	Use the getsensorinfo command. For power supplies, you can also use the System.Power.Supply command with the get subcommand. For more information, see the <i>iDRAC</i>
		www.dell.com/idracmanuals.
Fan	Dashboard > > System Health > Fans	
CPU	Dashboard > System Health > CPU	



Table 17. Sensor information using web interface and RACADM (continued)

View sensor information For	Using web interface	Using RACADM
Memory	Dashboard > System Health > Memory	
Intrusion	Dashboard > System Health > Intrusion	
Power Supplies	> Hardware > Power Supplies	
Removable Flash Media	Dashboard > System Health > Removable Flash Media	
Temperature	Dashboard > System Health > Power/Thermal > Temperatures	
Voltage	Dashboard > System Health > Power/Thermal > Voltages	

Monitoring performance index of CPU, memory, and input output modules

In Dell's 14th generation Dell PowerEdge servers, Intel ME supports Compute Usage Per Second (CUPS) functionality. The CUPS functionality provides real-time monitoring of CPU, memory, and I/O utilization and system-level utilization index for the system. Intel ME allows out-of-band (OOB) performance monitoring and does not consume CPU resources. The Intel ME has a system CUPS sensor that provides computation, memory, and I/O resource utilization values as a CUPS Index. iDRAC monitors this CUPS index for the overall system utilization and also monitors the instantaneous utilization index of the CPU, Memory, and I/O.

(i) **NOTE:** CUPS functionality is not supported on following servers:

- PowerEdge R240
- PowerEdge R240xd
- PowerEdge R340
- PowerEdge R6415
- PowerEdge R7415
- PowerEdge R7425
- PowerEdge T140

The CPU and chipset have dedicated Resource monitoring Counters (RMC). The data from these RMCs is queried to obtain utilization information of system resources. The data from RMCs is aggregated by the node manager to measure the cumulative utilization of each of these system resources that is read from iDRAC using existing intercommunication mechanisms to provide data through out-of-band management interfaces.

The Intel sensor representation of performance parameters and index values is for complete physical system. Therefore, the performance data representation on the interfaces is for the complete physical system, even if the system is virtualized and has multiple virtual hosts.

To display the performance parameters, the supported sensors must be present in the server.

The four system utilization parameters are:

- **CPU Utilization** Data from RMCs for each CPU core is aggregated to provide cumulative utilization of all the cores in the system. This utilization is based on time spent in active and inactive states. A sample of RMC is taken every six seconds.
- Memory Utilization RMCs measure memory traffic occurring at each memory channel or memory controller instance. Data from these RMCs is aggregated to measure the cumulative memory traffic across all the memory channels on the system. This is a measure of memory bandwidth consumption and not amount of memory utilization. iDRAC aggregates it for one minute, so it may or may not match the memory utilization that other OS tools, such as **top** in Linux, show. Memory bandwidth utilization that the iDRAC shows is an indication of whether workload is memory intensive or not.
- I/O Utilization There is one RMC per root port in the PCI Express Root Complex to measure PCI Express traffic emanating from or directed to that root port and the lower segment. Data from these RMCs is aggregated for measuring PCI express traffic for all PCI Express segments emanating from the package. This is measure of I/O bandwidth utilization for the system.

Viewing iDRAC and managed system information 117



• System Level CUPS Index — The CUPS index is calculated by aggregating CPU, Memory, and I/O index considering a predefined load factor of each system resource. The load factor depends on the nature of the workload on the system. CUPS Index represents the measurement of the compute headroom available on the server. If the system has a large CUPS Index, then there is limited headroom to place more workload on that system. As the resource consumption decreases, the system's CUPS index decreases. A low CUPS index indicates that there is a large compute headroom and the server can receive new workloads and the server is in a lower power state to reduce power consumption. Workload monitoring can then be applied throughout the data center to provide a high-level and holistic view of the data center's workload, providing a dynamic data center solution.

() NOTE: The CPU, memory, and I/O utilization indexes are aggregated over one minute. Therefore, if there are any

instantaneous spikes in these indexes, they may be suppressed. They are indication of workload patterns not the amount of resource utilization.

The IPMI, SEL, and SNMP traps are generated if the thresholds of the utilization indexes are reached and the sensor events are enabled. The sensor event flags are disabled by default. It can be enabled using the standard IPMI interface.

The required privileges are:

- Login privilege is required to monitor performance data.
- Configure privilege is required for setting warning thresholds and reset historical peaks.
- Login privilege and Enterprise license are required to read historical statics data.

Monitoring performance index of CPU, memory, and input output modules using web interface

To monitor the performance index of CPU, memory, and I/O modules, in the iDRAC web interface, go to **System** > **Performance**.

- System Performance section Displays the current reading and the warning reading for CPU, Memory and I/O utilization index, and system level CUPS index in a graphical view.
- System Performance Historical Data section:
 - Provides the statistics for CPU, memory, IO utilization, and the system level CUPS index. If the host system is powered off, then the graph displays the power off line below 0 percent.
 - You can reset the peak utilization for a particular sensor. Click **Reset Historical Peak**. You must have Configure privilege to reset the peak value.
- Performance Metrics section:
 - Displays status and present reading
 - Displays or specifies the warning threshold utilization limit. You must have server configure privilege to set the threshold values.

For information about the displayed properties, see the *iDRAC Online Help*.

Monitoring performance index for of CPU, memory, and input output modules using RACADM

Use the **SystemPerfStatistics** sub command to monitor performance index for CPU, memory, and I/O modules. For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Idle Server Detection

iDRAC provides out-of-band performance monitoring index of server components like CPU, memory, and I/O.

The history data of the server level CUPS index is used to monitor whether the server is utilized or running idle for long time. If the server is underutilized below certain threshold for a defined span of interval (in hours), then it will be reported as idle server.

This feature is only supported on Intel platforms with CUPS ability. AMD and Intel platforms without CUPS capability do not support this feature.

() NOTE:

- This feature requires Datacenter license.
- 118 Viewing iDRAC and managed system information



 To read the configurations of Idle Server Configuration parameters, you need Login privilege and to modify the parameters you need iDRAC Configure privilege.

To view or modify the parameters, navigate to **Configuration** > **System Settings**.

Idle server detection is reported based on following parameters:

- Idle Server Threshold (%) This is set to 20% by default and can be configured from 0 to 50%. The reset operation sets the threshold to 20%.
- Idle Server Scan Interval (in hours) This is the time period over which the hourly samples are collected to determine the idle server. This is set to 240 hours by default and can be configured from 1 to 9000 hours. The reset operation sets the interval to 240 hours.
- Server Utilization Percentile (%) The utilization percentile value can be set to 80 to 100%. The default value is 80%. If the 80% of the hourly samples falls below utilization threshold, then it is considered as idle server.

Modifying idle Server Detection parameters using RACADM

racadm get system.idleServerDetection

Modifying idle Server Detection parameters using Redfish

https://<iDRAC IP>/redfish/v1/Managers/System.Embedded.1/Attributes

Modifying idle Server Detection parameters using WSMAN

winrm e http://schemas.dmtf.org/wbem/wscim/1/cim-schema/2/root/dcim/DCIM_SystemAttribute -u:root -p:calvin -r:https://<iDRAC IP>/wsman -SkipCNcheck -SkipCAcheck -encoding:utf-8 -a:basic

i NOTE: iDRAC GUI doesn't support to view or modify the attributes.

GPU (Accelerators) Management

Dell PowerEdge servers are shipped with Graphics Processing Unit (GPU). GPU management enables you to view the various GPUs connected to the system and also monitor power, temperature, and thermal information for the GPUs.

NOTE: This is a licensed feature and is available only with iDRAC Datacenter and Enterprise licenses. Below properties require Datacenter/Enterprise license, other properties are listed even without these license:

- Thermal Metrics:
 - GPU Target Temperature
 - Minimum GPU HW Slowdown Temperature
 - GPU Shutdown Temperature
 - Maximum Memory Operating temperature
 - Maximum GPU Operating Temperature
 - Thermal Alert State
 - Power Brake State
- Power Metrics:
 - Power Supply Status
 - Board Power Supply Status
- Telemetry All GPU telemetry reports data

(i) NOTE: GPU properties will not be listed for Embedded GPU cards and the Status is marked as **Unknown**.



GPU has to be in ready state before the command fetches the data. GPUStatus field in Inventory shows the availability of the GPU and whether GPU device is responding or not. If the GPU status is ready, GPUStatus shows OK, otherwise the status shows Unavailable.

The GPU offers multiple health parameters which can be pulled through the SMBPB interface of the NVIDIA controllers. This feature is limited only to NVIDIA cards. Following are the health parameters retrieved from the GPU device:

- Power
- Temperature
- Thermal

NOTE: This feature is only limited to NVIDIA cards. This information is not available for any other GPU that the server may support. The interval for polling the GPU cards over the PBI is 5 seconds.

The host system must have the NVIDIA driver installed and running for the Power consumption, GPU target temperature, Min GPU slowdown temperature, GPU shutdown temperature, Max memory operating temperature, and Max GPU operating temperature features to be available. These values are shown as **N/A** if the GPU driver is not installed.

In Linux, when the card is unused, the driver down-trains the card and unloads in order to save power. In such cases, the Power consumption, GPU target temperature, Min GPU slowdown temperature, GPU shutdown temperature, Max memory operating temperature, and Max GPU operating temperature features are not available. Persistent mode should be enabled for the device to avoid unload. You can use nvidia-smi tool to enable this using the command nvidia-smi -pm 1.

You can generate GPU reports using Telemetry. For more information on telemetry feature, see Telemetry Streaming

NOTE: In Racadm, You may see dummy GPU entries with empty values. This may happen if device is not ready to respond when iDRAC queries the GPU device for the information. Perform iDRAC racrest operation to resolve this issue.

FPGA Monitoring

Field-programmable Gate Array (FPGA) devices needs real-time temperature sensor monitoring as it generates significant heat when in use. Perform the following steps to get FPGA inventory information:

- Power off the server.
- Install FPGA device on the riser card.
- Power on the server.
- Wait until POST is complete.
- Login to iDRAC GUI.
- Navigate to System > Overview > Accelerators. You can see both GPU and FPGA sections.
- Expand the specific FPGA component to see the following sensor information:
- Power consumption
- Temperature details

(i) NOTE: You must have iDRAC Login privilege to access FPGA information.

NOTE: Power consumption sensors are available only for the supported FPGA cards and is available only with Datacenter license.

Checking the system for Fresh Air compliance

Fresh Air cooling directly uses outside air to cool systems in the data center. Fresh Air compliant systems can operate above its normal ambient operating range (temperatures up to 113 °F (45 °C)).

NOTE: Some servers or certain configurations of a server may not be Fresh Air compliant. See the specific server manual for details related to Fresh Air compliance or contact Dell for more details.

To check the system for Fresh Air compliance:

- In the iDRAC Web interface, go to System > Overview > Cooling > Temperature overview. The Temperature overview page is displayed.
- 2. See the Fresh Air section that indicates whether the server is fresh air compliant or not.



Viewing historical temperature data

You can monitor the percentage of time the system has operated at ambient temperature that is greater than the normally supported fresh air temperature threshold. The system board temperature sensor reading is collected over a period of time to monitor the temperature. The data collection starts when the system is first powered on after it is shipped from the factory. The data is collected and displayed for the duration when the system is powered on. You can track and store the monitored temperature for the last seven years.

NOTE: You can track the temperature history even for systems that are not Fresh-Air compliant. However, the threshold limits and fresh air related warnings generated are based on fresh air supported limits. The limits are 42°C for warning and 47°C for critical. These values correspond to 40°C and 45°C fresh air limits with 2°C margin for accuracy.

Two fixed temperature bands are tracked that are associated to fresh air limits:

- Warning band Consists of the duration a system has operated above the temperature sensor warning threshold (42°C). The system can operate in the warning band for 10% of the time for 12 months.
- Critical band Consists of the duration a system has operated above the temperature sensor critical threshold (47°C). The system can operate in the critical band for 1% of the time for 12 months which also increments time in the warning band.

The collected data is represented in a graphical format to track the 10% and 1% levels. The logged temperature data can be cleared only before shipping from the factory.

An event is generated if the system continues to operate above the normally supported temperature threshold for a specified operational time. If the average temperature over the specified operational time is greater than or equal to the warning level (> = 8%) or the critical level (> = 0.8%), an event is logged in the Lifecycle Log and the corresponding SNMP trap is generated. The events are:

- Warning event when the temperature was greater than the warning threshold for duration of 8% or more in the last 12 months.
- Critical event when the temperature was greater than the warning threshold for duration of 10% or more in the last 12 months.
- Warning event when the temperature was greater than the critical threshold for duration of 0.8% or more in the last 12 months.
- Critical event when the temperature was greater than the critical threshold for duration of 1% or more in the last 12 months.

You can also configure iDRAC to generate additional events. For more information, see the Setting alert recurrence event section.

Viewing historical temperature data using iDRAC web interface

To view historical temperature data:

- In the iDRAC Web interface, go to System > Overview > Cooling > Temperature overview. The Temperature overview page is displayed.
- 2. See the **System Board Temperature Historical Data** section that provides a graphical display of the stored temperature (average and peak values) for the last day, last 30 days, and last year.

For more information, see the *iDRAC Online Help*.

(i) NOTE: After an iDRAC firmware update or iDRAC reset, some temperature data may not be displayed in the graph.

() NOTE: WX3200 AMD GPU card currently doesnot support I2C interface for temperature sensors. Hence, temperature readings will not be available for this card from iDRAC interfaces.

Viewing historical temperature data using RACADM

To view historical data using RACADM, use the inlettemphistory command.

For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.



Configuring warning threshold for inlet temperature

You can modify the minimum and maximum warning threshold values for the system board inlet temperature sensor. If reset to default action is performed, the temperature thresholds are set to the default values. You must have Configure user privilege to set the warning threshold values for the inlet temperature sensor.

Configuring warning threshold for inlet temperature using web interface

To configure warning threshold for inlet temperature:

- 1. In the iDRAC Web interface, go to **System** > **Overview** > **Cooling** > **Temperature overview**. The **Temperature overview** page is displayed.
- 2. In the Temperature Probes section, for the System Board Inlet Temp, enter the minimum and maximum values for the Warning Threshold in Centigrade or Fahrenheit. If you enter the value in centigrade, the system automatically calculates and displays the Fahrenheit value. Similarly, if you enter Fahrenheit, the value for Centigrade is displayed.
- 3. Click Apply.

The values are configured.

() NOTE: Changes to default thresholds are not reflected in the historical data chart since the chart limits are for fresh air limit values only. Warnings for exceeding the custom thresholds are different from warning associated to exceeding fresh air thresholds.

Viewing network interfaces available on host OS

You can view information about all the network interfaces that are available on the host operating system such as the IP addresses that are assigned to the server. The iDRAC Service Module provides this information to iDRAC. The OS IP address information includes the IPv4 and IPv6 addresses, MAC address, Subnet mask or prefix length, the FQDD of the network device, network interface name, network interface description, network interface status, network interface type (Ethernet, tunnel, loopback, and so on.), Gateway address, DNS server address, and DHCP server address.

(i) NOTE: This feature is available with iDRAC Express and iDRAC Enterprise/Datacenter licenses.

To view the OS information, make sure that:

- You have Login privilege.
- iDRAC Service Module is installed and running on the host operating system.
- OS Information option is enabled in the **iDRAC Settings** > **Overview** > **iDRAC Service Module** page.

iDRAC can display the IPv4 and IPv6 addresses for all the interfaces configured on the Host OS.

Depending on how the Host OS detects the DHCP server, the corresponding IPv4 or IPv6 DHCP server address may not be displayed.

Viewing network interfaces available on host OS using web interface

To view the network interfaces available on the host OS using Web interface:

- 1. Go to System > Host OS > Network Interfaces.
- The Network Interfaces page displays all the network interfaces that are available on the host operating system.
- To view the list of network interfaces associated with a network device, from the Network Device FQDD drop-down menu, select a network device and click Apply. The OS IP details are displayed in the Host OS Network Interfaces section.
- The US IP details are displayed in the **Host US Network Interfaces** sect
- 3. From the Device FQDD column, click on the network device link. The corresponding device page is displayed from the Hardware > Network Devices section, where you can view the device details. For information about the properties, see the *iDRAC Online Help*.
- 4. Click the 🛨 icon to display more details.

Similarly, you can view the host OS network interface information associated with a network device from the **Hardware** > **Network Devices** page. Click **View Host OS Network Interfaces**.

122 Viewing iDRAC and managed system information



NOTE: For the ESXi host OS in the iDRAC Service Module v2.3.0 or later, the **Description** column in the **Additional Details** list is displayed in the following format:

<List-of-Uplinks-Configured-on-the-vSwitch>/<Port-Group>/<Interface-name>

Viewing network interfaces available on host OS using RACADM

Use the gethostnetworkinterfaces command to view the network interfaces available on the host operating systems using RACADM. For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Viewing FlexAddress mezzanine card fabric connections

In blade servers, FlexAddress allows the use of persistent, chassis-assigned World Wide Names and MAC addresses (WWN/MAC) for each managed server port connection.

You can view the following information for each installed embedded Ethernet and optional mezzanine card port:

- Fabrics to which the cards are connected.
- Type of fabric.
- Server-assigned, chassis-assigned, or remotely assigned MAC addresses.

To view the Flex Address information in iDRAC, configure and enable the Flex Address feature in Chassis Management Controller (CMC). For more information, see the *Chassis Management Controller User's Guide* available at https://www.dell.com/cmcmanuals. Any existing Virtual Console or Virtual Media session terminates if the FlexAddress setting is enabled or disabled.

NOTE: To avoid errors that may lead to an inability to turn on the managed system, you *must* have the correct type of mezzanine card installed for each port and fabric connection.

The FlexAddress feature replaces the server–assigned MAC addresses with chassis–assigned MAC addresses and is implemented for iDRAC along with blade LOMs, mezzanine cards and I/O modules. The iDRAC FlexAddress feature supports preservation of slot specific MAC address for iDRACs in a chassis. The chassis–assigned MAC address is stored in CMC non–volatile memory and is sent to iDRAC during an iDRAC boot or when CMC FlexAddress is enabled.

If CMC enables chassis-assigned MAC addresses, iDRAC displays the MAC address on any of the following pages:

- System > Details > iDRAC Details.
- System > Server > WWN/MAC.
- iDRAC Settings > Overview > Current Network Settings.

CAUTION: With FlexAddress enabled, if you switch from a server-assigned MAC address to a chassis-assigned MAC address and vice-versa, iDRAC IP address also changes.

Viewing or terminating iDRAC sessions

You can view the number of users currently logged in to iDRAC and terminate the user sessions.

Terminating iDRAC sessions using web interface

The users who do not have administrative privileges must have Configure iDRAC privilege to terminate iDRAC sessions using iDRAC Web interface.

To view and terminate the iDRAC sessions:

1. In the iDRAC Web interface, go to **iDRAC Settings** > **users** > **Sessions**.



The **Sessions** page displays the session ID, username, IP address, and session type. For more information about these properties, see the *iDRAC Online Help*.

2. To terminate the session, under the **Terminate** column, click the Trashcan icon for a session.

Terminating iDRAC sessions using RACADM

You must have administrator privileges to terminate iDRAC sessions using RACADM.

To view the current user sessions, use the getssninfo command.

To terminate a user session, use the closessn command.

For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.



Setting up iDRAC communication

You can communicate with iDRAC using any of the following modes:

- iDRAC Web Interface
- Serial connection using DB9 cable (RAC serial or IPMI serial) For rack and tower servers only
- IPMI Serial Over LAN
- IPMI Over LAN
- Remote RACADM
- Local RACADM
- Remote Services
- **NOTE:** To ensure that Local RACADM import or export commands work properly, ensure that the USB mass-storage host is enabled in the operating system. For information about enabling USB storage host, see the documentation for your operating system.

The following table provides an overview of the supported protocols, supported commands, and pre-requisites:

Table 18. Communication modes — summary

Mode of Communication	Supported Protocol	Supported Commands	Pre-requisite
iDRAC Web Interface	Internet Protocol (https)	N/A	Web Server
Serial using Null modem	Serial Protocol	RACADM	Part of iDRAC firmware
		IPMI	RAC Serial or IPMI Serial is enabled
IPMI Serial Over LAN	Intelligent Platform Management Bus protocol SSH	IPMI	IPMITool is installed and IPMI Serial Over LAN is enabled
IPMI over LAN	Intelligent Platform Management Bus protocol	IPMI	IPMITool is installed and IPMI Settings is enabled
Remote RACADM	https	Remote RACADM	Remote RACADM is installed and enabled
Firmware RACADM	SSH	Firmware RACADM	Firmware RACADM is installed and enabled
Local RACADM	IPMI	Local RACADM	Local RACADM is installed
Remote Services ¹	WSMan	WinRM (Windows)	WinRM is installed (Windows)
		OpenWSMan (Linux)	(Linux)
	Redfish	Various browser plug-ins, CURL (Windows and Linux), Python request and JSON modules	Plug-ins, CURL, Python modules are installed
[1] For more information, see	the Lifecycle Controller User's (Juide available at https://www.de	Il.com/idracmanuals.

Topics:

- Communicating with iDRAC through serial connection using DB9 cable
- Switching between RAC serial and serial console while using DB9 cable
- Communicating with iDRAC using IPMI SOL
- Communicating with iDRAC using IPMI over LAN

Inserido ao protocolo 22.951.206-4 por: Pedro Henrique Golin Linhares em: 31/03/2025 12:49. A autenticidade deste documento pode ser validada no endereço: https://www.eprotocolo.pr.gov.br/spiweb/validarDocumento com o código: 692a5ef82beb173ed33dad4ee6505546.



- Enabling or disabling remote RACADM
- Disabling local RACADM
- Enabling IPMI on managed system
- Configuring Linux for serial console during boot in RHEL 6
- Configuring serial terminal in RHEL 7
- Supported SSH cryptography schemes

Communicating with iDRAC through serial connection using DB9 cable

You can use any of the following communication methods to perform systems management tasks through serial connection to rack and tower servers:

- RAC Serial
- IPMI Serial Direct Connect Basic mode and Direct Connect Terminal mode
- **NOTE:** In case of blade servers, the serial connection is established through the chassis. For more information, see the *Chassis Management Controller User's Guide* available at https://www.dell.com/cmcmanuals (not applicable for MX platforms) *OME Modular for PowerEdge MX7000 Chassis User's Guide* available at https://www.dell.com/cmcmanuals (not applicable of MX7000 Chassis User's Guide available at https://www.dell.com/cmcmanuals (not applicable of MX7000 Chassis User's Guide available at https://www.dell.com/openmanagemanuals (applicable for MX platforms).

To establish the serial connection:

- **1.** Configure the BIOS to enable serial connection.
- 2. Connect the Null Modem DB9 cable from the management station's serial port to the managed system's external serial connector.

(i) NOTE: Server power cycle is required from vConsole or GUI for any change in Baud-rate.

NOTE: If iDRAC serial connection authentication is disabled, then iDRAC racreset is required for any change in BAUD-rate.

- **3.** Make sure that the management station's terminal emulation software is configured for serial connection using any of the following:
 - Linux Minicom in an Xterm
 - Hilgraeve's HyperTerminal Private Edition (version 6.3)

Based on where the managed system is in its boot process, you can see either the POST screen or the operating system screen. This is based on the configuration: SAC for Windows and Linux text mode screens for Linux.

4. Enable RAC serial or IPMI serial connections in iDRAC.

Configuring BIOS for serial connection

To configure BIOS for Serial Connection:

(i) NOTE: This is applicable only for iDRAC on rack and tower servers.

- 1. Turn on or restart the system.
- 2. Press F2.
- 3. Go to System BIOS Settings > Serial Communication.
- 4. Select External Serial Connector to Remote Access device.
- 5. Click Back, click Finish, and then click Yes.
- 6. Press Esc to exit System Setup.

Enabling RAC serial connection

After configuring serial connection in BIOS, enable RAC serial in iDRAC.

126 Setting up iDRAC communication



(i) NOTE: This is applicable only for iDRAC on rack and tower servers.

Enabling RAC serial connection using web interface

To enable RAC serial connection:

- In the iDRAC Web interface, go to iDRAC Settings > Network > Serial. The Serial page is displayed.
- 2. Under RAC Serial, select Enabled and specify the values for the attributes.
- 3. Click Apply.
 - The RAC serial settings are configured.

Enabling RAC serial connection using RACADM

To enable RAC serial connection using RACADM, use the set command with the object in the iDRAC. Serial group.

Enabling IPMI serial connection basic and terminal modes

To enable IPMI serial routing of BIOS to iDRAC, configure IPMI Serial in any of the following modes in iDRAC:

(i) NOTE: This is applicable only for iDRAC on rack and tower servers.

 IPMI basic mode — Supports a binary interface for program access, such as the IPMI shell (ipmish) that is included with the Baseboard Management Utility (BMU). For example, to print the System Event Log using ipmish via IPMI Basic mode, run the following command:

ipmish -com 1 -baud 57600 -flow cts -u <username> -p <password> sel get

(i) NOTE: The default iDRAC user name and password are provided on the system badge.

- IPMI terminal mode Supports ASCII commands that are sent from a serial terminal. This mode supports limited number of commands (including power control) and raw IPMI commands that are typed as hexadecimal ASCII characters. It allows you to view the operating system boot sequences up to BIOS, when you login to iDRAC through SSH. You need to logout from the IPMI terminal using [sys pwd -x], below are the example for IPMI Terminal mode commands.
 - o [sys tmode]

```
o [sys pwd -u root calvin]
```

- o [sys health query -v]
- 0 [18 00 01]
- [sys pwd -x]

Enabling serial connection using web interface

Make sure to disable the RAC serial interface to enable IPMI Serial.

To configure IPMI Serial settings:

- 1. In the iDRAC Web interface, go to iDRAC Settings > Connectivity > Serial.
- 2. Under IPMI Serial, specify the values for the attributes. For information about the options, see the iDRAC Online Help.
- 3. Click Apply.

Enabling serial connection IPMI mode using RACADM

To configure the IPMI mode, disable the RAC serial interface and then enable the IPMI mode.

```
racadm set iDRAC.Serial.Enable 0
racadm set iDRAC.IPMISerial.ConnectionMode <n>
```

n=0 — Terminal Mode



n=1 — Basic Mode

Enabling serial connection IPMI serial settings using RACADM

1. Change the IPMI serial-connection mode to the appropriate setting using the command.

racadm set iDRAC.Serial.Enable 0

2. Set the IPMI Serial baud rate using the command.

```
racadm set iDRAC.IPMISerial.BaudRate <baud rate>
```

Parameter	Allowed values (in bps)
<baud_rate></baud_rate>	9600, 19200, 57600, and 115200.

3. Enable the IPMI serial hardware flow control using the command.

racadm set iDRAC.IPMISerial.FlowContro 1

4. Set the IPMI serial channel minimum privilege level using the command.

racadm set iDRAC.IPMISerial.ChanPrivLimit <level>

Parameter	Privilege level
<level> = 2</level>	User
<level> = 3</level>	Operator
<level> = 4</level>	Administrator

5. Ensure that the serial MUX (external serial connector) is set correctly to the remote access device in the BIOS Setup program to configure BIOS for serial connection.

For more information about these properties, see the IPMI 2.0 specification.

Additional settings for ipmi serial terminal mode

This section provides additional configuration settings for IPMI serial terminal mode.

Configuring additional settings for IPMI serial terminal mode using web interface

To set the Terminal Mode settings:

- In the iDRAC Web interface, go to iDRAC Settings > Connectivity > Serial. The Serial page is displayed.
- 2. Enable IPMI serial.
- **3.** Click **Terminal Mode Settings.** The **Terminal Mode Settings** page is displayed.
- **4.** Specify the following values:
 - Line editing
 - Delete control
 - Echo Control
 - Handshaking control
 - New line sequence
 - Input new line sequences

For information about the options, see the *iDRAC Online Help*.



5. Click Apply.

The terminal mode settings are configured.

6. Make sure that the serial MUX (external serial connector) is set correctly to the remote access device in the BIOS Setup program to configure BIOS for serial connection.

Configuring additional settings for IPMI serial terminal mode using RACADM

To configure the Terminal Mode settings, use the set command with the objects in the idrac.ipmiserial group. For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Switching between RAC serial and serial console while using DB9 cable

iDRAC supports Escape key sequences that allow switching between RAC Serial Interface communication and Serial Console on rack and tower servers.

Switching from serial console to RAC serial

To switch to RAC Serial Interface communication mode when in Serial Console Mode, press Esc+Shift, 9.

The key sequence directs you to the iDRAC Login prompt (if the iDRAC is set to RAC Serial mode) or to the Serial Connection mode where terminal commands can be issued if iDRAC is set to IPMI Serial Direct Connect Terminal Mode.

Switching from RAC serial to serial console

To switch to Serial Console Mode when in RAC Serial Interface Communication Mode, press Esc+Shift, Q.

When in terminal mode, to switch the connection to the Serial Console mode, press Esc+Shift, Q.

To go back to the terminal mode use, when connected in Serial Console mode, press Esc+Shift, 9.

Communicating with iDRAC using IPMI SOL

IPMI Serial Over LAN (SOL) allows a managed system's text-based console serial data to be redirected over iDRAC's dedicated or shared out-of-band ethernet management network. Using SOL you can:

- Remotely access operating systems with no time-out.
- Diagnose host systems on Emergency Management Services (EMS) or Special Administrator Console (SAC) for Windows or Linux shell.
- View the progress of a servers during POST and reconfigure the BIOS setup program.
- To setup the SOL communication mode:
- 1. Configure BIOS for serial connection.
- **2.** Configure iDRAC to Use SOL.
- 3. Enable a supported protocol (SSH, IPMItool).

Configuring BIOS for serial connection

(i) NOTE: This is applicable only for iDRAC on rack and tower servers.

- 1. Turn on or restart the system.
- 2. Press F2.
- 3. Go to System BIOS Settings > Serial Communication.
- **4.** Specify the following values:



- Serial Communication On With Console Redirection
- Serial Port Address COM2.
 NOTE: You can set the serial communication field to On with serial redirection via com1 if serial device2 in the serial port address field is also set to com1.
- External serial connector Serial device 2
- Failsafe Baud Rate 115200
- Remote Terminal Type VT100/VT220
- Redirection After Boot Enabled
- 5. Click **Back** and then click **Finish**.
- 6. Click Yes to save the changes.
- 7. Press <Esc> to exit System Setup.
 - **NOTE:** BIOS sends the screen serial data in 25 x 80 format. The SSH window that is used to invoke the console com2 command must be set to 25 x 80. Then, the redirected screen appears correctly.
 - () NOTE: If the boot loader or operating system provides serial redirection such as GRUB or Linux, then the BIOS **Redirection After Boot** setting must be disabled. This is to avoid potential race condition of multiple components accessing the serial port.

Configuring iDRAC to use SOL

You can specify the SOL settings in iDRAC using Web interface, RACADM, or iDRAC Settings utility.

Configuring iDRAC to use SOL using iDRAC web interface

To configure IPMI Serial over LAN (SOL):

- In the iDRAC Web interface, go to iDRAC Settings > Connectivity > Serial Over LAN. The Serial over LAN page is displayed.
- **2.** Enable SOL, specify the values, and click **Apply**. The IPMI SOL settings are configured.
- **3.** To set the character accumulate interval and the character send threshold, select **Advanced Settings**. The **Serial Over LAN Advanced Settings** page is displayed.
- 4. Specify the values for the attributes and click Apply.

The IPMI SOL advanced settings are configured. These values help to improve the performance.

For information about the options, see the *iDRAC Online Help*.

Configuring iDRAC to use SOL using RACADM

To configure IPMI Serial over LAN (SOL):

1. Enable IPMI Serial over LAN using the command.

racadm set iDRAC.IPMISol.Enable 1

2. Update the IPMI SOL minimum privilege level using the command.

racadm set iDRAC.IPMISol.MinPrivilege <level>

Parameter	Privilege level
<level> = 2</level>	User
<level> = 3</level>	Operator
<level> = 4</level>	Administrator



NOTE: To activate IPMI SOL, you must have the minimum privilege defined in IMPI SOL. For more information, see the IPMI 2.0 specification.

3. Update the IPMI SOL baud rate using the command.

racadm set iDRAC.IPMISol.BaudRate <baud_rate>

NOTE: To redirect the serial console over LAN, make sure that the SOL baud rate is identical to the managed system's baud rate.

Parameter	Allowed values (in bps)
<baud_rate></baud_rate>	9600, 19200, 57600, and 115200.

4. Enable SOL for each user using the command.

racadm set iDRAC.Users.<id>.SolEnable 2

Parameter	Description
<id></id>	Unique ID of the user

NOTE: To redirect the serial console over LAN, ensure that the SOL baud rate is identical to the baud rate of the managed system.

Enabling supported protocol

The supported protocols are IPMI and SSH.

Enabling supported protocol using web interface

To enable SSH, go to **iDRAC Settings** > Services and select Enabled for SSH.

To enable IPMI, go to **iDRAC Settings** > **Connectivity** and select **IPMI Settings**. Make sure that the **Encryption Key** value is all zeroes or press the backspace key to clear and change the value to NULL characters.

Enabling supported protocol using RACADM

To enable the SSH, use the following command.

SSH

```
racadm set iDRAC.SSH.Enable 1
```

To change the SSH port

racadm set iDRAC.SSH.Port <port number>

You can use tools such as:

- IPMItool for using IPMI protocol
- Putty/OpenSSH for using SSH protocol

SOL using IPMI protocol

The IPMI-based SOL utility and IPMItool use RMCP+ delivered using UDP datagrams to port 623. The RMCP+ provides improved authentication, data integrity checks, encryption, and the ability to carry multiple types of payloads while using IPMI 2.0. For more information, see http://ipmitool.sourceforge.net/manpage.html.

Setting up iDRAC communication 131



The RMCP+ uses a 40-character hexadecimal string (characters 0-9, a-f, and A-F) encryption key for authentication. The default value is a string of 40 zeros.

An RMCP+ connection to iDRAC must be encrypted using the encryption key (Key Generator Key). You can configure the encryption key using the iDRAC web interface or iDRAC Settings utility.

To start SOL session using IPMItool from a management station:

(i) NOTE: If required, you can change the default SOL time-out at iDRAC Settings > Services.

- 1. Install IPMITool from the *Dell Systems Management Tools and Documentation* DVD. For installation instructions, see the *Software Quick Installation Guide*.
- 2. At the command prompt (Windows or Linux), run the following command to start SOL from iDRAC:

```
ipmitool -H <iDRAC-ip-address> -I lanplus -U <login name> -P <login password> sol activate
```

This command connected the management station to the managed system's serial port.

3. To quit a SOL session from IPMItool, press \sim and then . (period).

(i) NOTE: If a SOL session does not terminate, reset iDRAC and allow up to two minutes to complete booting.

- **NOTE:** IPMI SOL session may terminate while copying large input text from a client running Windows OS to a host running Linux OS. To avoid the session from getting terminated abruptly, convert any large text to a UNIX-based line ending.
- **NOTE:** If a SOL session created using RACADM tool exists, starting another SOL session using IPMI tool will not show any notification or error about the existing sessions.
- **NOTE:** Due to windows OS settings, SOL session connected through ssh and IPMI tool may go to blank screen after booting. Disconnect and Re-connect the SOL Session again to get back SAC prompt.

SOL using SSH

Secure Shell (SSH) is network protocol used to perform command line communications to iDRAC. You can parse remote RACADM commands through this interface.

SSH has improved security. iDRAC only supports SSH version 2 with password authentication, and is enabled by default. iDRAC supports up to two to four SSH sessions at a time.

- () NOTE: Starting iDRAC version 4.40.00.00, telnet feature is removed from iDRAC, so any related attribute registry properties are obsoleted. While some of these properties are still available in iDRAC to keep backward compatibility with existing console applications and scripts, the corresponding settings are ignored by iDRAC firmware.
- (i) NOTE: While establishing SSH connection, a security message is displayed- 'Further Authentication required'. even though 2FA is disabled.
- **NOTE:** For MX platforms, one SSH session will be used for iDRAC communication. If all the sessions are in use, then iDRAC will not launch until one session is free.

Use open-source programs such as PuTTY or OpenSSH that support SSH on a management station to connect to iDRAC.

NOTE: Run OpenSSH from a VT100 or ANSI terminal emulator on Windows. Running OpenSSH at the Windows command prompt does not result in full functionality (that is, some keys do not respond and no graphics are displayed).

Before using SSH to communicate with iDRAC, make sure to:

- 1. Configure BIOS to enable Serial Console.
- 2. Configure SOL in iDRAC.
- 3. Enable SSH using iDRAC Web interface or RACADM.
 - SSH (port 22) client <--> WAN connection <--> iDRAC

The IPMI-based SOL that uses SSH protocol eliminates the need for an additional utility because the serial to network translation happens within iDRAC. The SSH console that you use must be able to interpret and respond to the data arriving



from the serial port of the managed system. The serial port usually attaches to a shell that emulates an ANSI- or VT100/VT220-terminal. The serial console is automatically redirected to the SSH.

Using SOL from PuTTY on Windows

(i) NOTE: If required, you can change the default SSH time-out at **iDRAC Settings** > Services.

To start IPMI SOL from PuTTY on a Windows management station:

1. Run the following command to connect to iDRAC

putty.exe [-ssh] <login name>@<iDRAC-ip-address> <port number>

(i) NOTE: The port number is optional. It is required only when the port number is reassigned.

2. Run the command console com2 or connect to start SOL and boot the managed system.

A SOL session from the management station to the managed system using the SSH protocol is opened. To access the iDRAC command-line console, follow the ESC key sequence. Putty and SOL connection behavior:

- While accessing the managed system through putty during POST, if the Function keys and keypad option on putty is set to:
 - VT100+ F2 passes, but F12 cannot pass.
 - ESC[n~ F12 passes, but F2 cannot pass.
- In Windows, if the Emergency Management System (EMS) console is opened immediately after a host reboot, the Special Admin Console (SAC) terminal may get corrupted. Quit the SOL session, close the terminal, open another terminal, and start the SOL session using the same command.
- **NOTE:** Due to windows OS settings, SOL session connected through ssh and IPMI tool may go to blank screen after booting. Disconnect and Re-connect the SOL Session again to get back SAC prompt.

Using SOL from OpenSSH on Linux

To start SOL from OpenSSH on a Linux management station:

(i) NOTE: If required, you can change the default SSH session time-out at iDRAC Settings > Services.

- 1. Start a shell.
- 2. Connect to iDRAC using the following command: ssh <iDRAC-ip-address> -I <login name>
- **3.** Enter one of the following commands at the command prompt to start SOL:
 - connect
 - console com2

This connects iDRAC to the managed system's SOL port. Once a SOL session is established, iDRAC command line console is not available. Follow the escape sequence correctly to open the iDRAC command line console. The escape sequence is also printed on the screen as soon as a SOL session is connected. When the managed system is off, it takes sometime to establish the SOL session.

(i) NOTE: You can use console com1 or console com2 to start SOL. Reboot the server to establish the connection.

The console -h com2 command displays the contents of the serial history buffer before waiting for input from the keyboard or new characters from the serial port.

The default (and maximum) size of the history buffer is 8192 characters. You can set this number to a smaller value using the command:

racadm set iDRAC.Serial.HistorySize <number>

4. Quit the SOL session to close an active SOL session.



Disconnecting SOL session in iDRAC command line console

The commands to disconnect a SOL session are based on the utility. You can exit the utility only when a SOL session is completely terminated.

To disconnect a SOL session, terminate the SOL session from the iDRAC command line console.

 To quit SOL redirection, press Enter, Esc, T. The SOL session closes.

If a SOL session is not terminated completely in the utility, other SOL sessions may not be available. To resolve this, terminate the command line console in the Web interface under **iDRAC Settings** > **Connectivity** > **Serial Over LAN**.

Communicating with iDRAC using IPMI over LAN

You must configure IPMI over LAN for iDRAC to enable or disable IPMI commands over LAN channels to any external systems. If IPMI over LAN is not configured, then external systems cannot communicate with the iDRAC server using IPMI commands.

(i) NOTE: IPMI also supports IPv6 address protocol for Linux-based operating systems.

Configuring IPMI over LAN using web interface

To configure IPMI over LAN:

- In the iDRAC Web interface, go to iDRAC Settings > Connectivity. The Network page is displayed.
- 2. Under IPMI Settings, specify the values for the attributes and click Apply.

For information about the options, see the *iDRAC Online Help*.

The IPMI over LAN settings are configured.

Configuring IPMI over LAN using iDRAC settings utility

To configure IPMI over LAN:

- In the iDRAC Settings Utility, go to Network. The iDRAC Settings Network page is displayed.
- 2. For IPMI Settings, specify the values.

For information about the options, see the *iDRAC Settings Utility Online Help*.

3. Click **Back**, click **Finish**, and then click **Yes**. The IPMI over LAN settings are configured.

Configuring IPMI over LAN using RACADM

1. Enable IPMI over LAN.

```
racadm set iDRAC.IPMILan.Enable 1
```

- (i) **NOTE:** This setting determines the IPMI commands that are executed using IPMI over LAN interface. For more information, see the IPMI 2.0 specifications at **intel.com**.
- 2. Update the IPMI channel privileges.

```
racadm set iDRAC.IPMILan.PrivLimit <level>
```

Parameter	Privilege level
<level> = 2</level>	User
<level> = 3</level>	Operator
<level> = 4</level>	Administrator

3. Set the IPMI LAN channel encryption key , if required.

racadm set iDRAC.IPMILan.EncryptionKey <key>

Parameter	Description
<key></key>	20-character encryption key in a valid hexadecimal format.

(i) NOTE: The iDRAC IPMI supports the RMCP+ protocol. For more information, see the IPMI 2.0 specifications at intel.com.

Enabling or disabling remote RACADM

You can enable or disable remote RACADM using the iDRAC Web interface or RACADM. You can run up to five remote RACADM sessions in parallel.

(i) NOTE: Remote RACADM is enabled by default.

Enabling or disabling remote RACADM using web interface

- 1. In iDRAC Web interface, go to **iDRAC Settings** > **Services**.
- Under Remote RACADM, select the desired option and click Apply. The remote RACADM is enabled or disabled based on the selection.

Enabling or disabling remote RACADM using RACADM

(i) NOTE: It is recommended to run these commands using local RACADM or firmware RACADM.

• To disable remote RACADM:

racadm set iDRAC.Racadm.Enable 0

To enable remote RACADM:

racadm set iDRAC.Racadm.Enable 1

Disabling local RACADM

The local RACADM is enabled by default. To disable, see Disabling access to modify iDRAC configuration settings on host system.

Enabling IPMI on managed system

On a managed system, use the Dell Open Manage Server Administrator to enable or disable IPMI. For more information, see the *OpenManage Server Administrator User's Guide* available at https://www.dell.com/openmanagemanuals.

(i) NOTE: From iDRAC v2.30.30.30 or later, IPMI supports IPv6 address protocol for Linux-based operating systems.

Setting up iDRAC communication 135

Inserido ao protocolo **22.951.206-4** por: **Pedro Henrique Golin Linhares** em: 31/03/2025 12:49. A autenticidade deste documento pode ser validada no endereço: https://www.eprotocolo.pr.gov.br/spiweb/validarDocumento com o código: 692a5ef82beb173ed33dad4ee6505546.



Configuring Linux for serial console during boot in RHEL 6

The following steps are specific to the Linux GRand Unified Bootloader (GRUB). Similar changes are required if a different boot loader is used.

NOTE: When you configure the client VT100 emulation window, set the window or application that is displaying the redirected Virtual Console to 25 rows x 80 columns to make sure the correct text displays. Else, some text screens may be garbled.

Edit the /etc/grub.conf file as follows:

1. Locate the General Setting sections in the file and add the following:

```
serial --unit=1 --speed=57600 terminal --timeout=10 serial
```

2. Append two options to the kernel line:

kernel console=ttyS1,115200n8r console=tty1

3. Disable GRUB's graphical interface and use the text-based interface. Else, the GRUB screen is not displayed in RAC Virtual Console. To disable the graphical interface, comment-out the line starting with splashimage.

The following example provides a sample /etc/grub.conf file that shows the changes described in this procedure.

```
# grub.conf generated by anaconda
 Note that you do not have to rerun grub after making changes to this file
#
# NOTICE: You do not have a /boot partition. This means that all
# kernel and initrd paths are relative to /, e.g.
#
 root (hd0,0)
 kernel /boot/vmlinuz-version ro root=/dev/sdal
# initrd /boot/initrd-version.img
#boot=/dev/sda
default=0
timeout=10
#splashimage=(hd0,2)/grub/splash.xpm.gz
serial --unit=1 --speed=57600
terminal --timeout=10 serial
title Red Hat Linux Advanced Server (2.4.9-e.3smp) root (hd0,0)
kernel /boot/vmlinuz-2.4.9-e.3smp ro root=/dev/sda1 hda=ide-scsi console=ttyS0
console=ttyS1,115200n8r
initrd /boot/initrd-2.4.9-e.3smp.img
title Red Hat Linux Advanced Server-up (2.4.9-e.3) root (hd0,00)
kernel /boot/vmlinuz-2.4.9-e.3 ro root=/dev/sda1 s
initrd /boot/initrd-2.4.9-e.3.im
```

4. To enable multiple GRUB options to start Virtual Console sessions through the RAC serial connection, add the following line to all options:

console=ttyS1,115200n8r console=tty1

The example shows console=ttyS1, 57600 added to the first option.

(i) NOTE: If the boot loader or operating system provides serial redirection such as GRUB or Linux, then the BIOS Redirection After Boot setting must be disabled. This is to avoid potential race condition of multiple components accessing the serial port.

Enabling login to the virtual console after boot

In the file /etc/inittab, add a new line to configure agetty on the COM2 serial port:


co:2345:respawn:/sbin/agetty -h -L 57600 ttyS1 ansi

The following example shows a sample file with the new line.

#inittab This file describes how the INIT process should set up #the system in a certain run-level. #Author:Miquel van Smoorenburg #Modified for RHS Linux by Marc Ewing and Donnie Barnes #Default runlevel. The runlevels used by RHS are: #0 - halt (Do NOT set initdefault to this) #1 - Single user mode #2 - Multiuser, without NFS (The same as 3, if you do not have #networking) #3 - Full multiuser mode #4 - unused #5 - X11 #6 - reboot (Do NOT set initdefault to this) id:3:initdefault: #System initialization. si::sysinit:/etc/rc.d/rc.sysinit 10:0:wait:/etc/rc.d/rc 0 l1:1:wait:/etc/rc.d/rc 1 l2:2:wait:/etc/rc.d/rc 13:3:wait:/etc/rc.d/rc 3 14:4:wait:/etc/rc.d/rc 4 15:5:wait:/etc/rc.d/rc 5 l6:6:wait:/etc/rc.d/rc 6 #Things to run in every runlevel. ud::once:/sbin/update ud::once:/sbin/update #Trap CTRL-ALT-DELETE ca::ctrlaltdel:/sbin/shutdown -t3 -r now #When our UPS tells us power has failed, assume we have a few #minutes of power left. Schedule a shutdown for 2 minutes from now. #This does, of course, assume you have power installed and your #UPS is connected and working correctly. pf::powerfail:/sbin/shutdown -f -h +2 "Power Failure; System Shutting Down" #If power was restored before the shutdown kicked in, cancel it. pr:12345:powerokwait:/sbin/shutdown -c "Power Restored; Shutdown Cancelled" #Run gettys in standard runlevels co:2345:respawn:/sbin/agetty -h -L 57600 ttyS1 ansi 1:2345:respawn:/sbin/mingetty tty1

2:2345:respawn:/sbin/mingetty tty2 3:2345:respawn:/sbin/mingetty tty3 4:2345:respawn:/sbin/mingetty tty4 5:2345:respawn:/sbin/mingetty tty5 6:2345:respawn:/sbin/mingetty tty6 #Run xdm in runlevel 5 #xdm is now a separate service x:5:respawn:/etc/X11/prefdm -nodaemon

In the file **/etc/securetty** add a new line with the name of the serial tty for COM2:

ttvS1

The following example shows a sample file with the new line.

NOTE: Use the Break Key Sequence (~B) to execute the Linux Magic SysRq key commands on serial console using IPMI Tool.

vc/1 vc/2 vc/3 vc/4 vc/5

vc/6

- vc/7 vc/8
- vc/9

Setting up iDRAC communication 137

	INTER	Fls. <u>524</u> Mov. <u>3</u>	$\frac{1}{2}$
vc/10 vc/11 tty1	35	ADO DO	Ê
tty2 tty3 tty4			
tty5 tty6 tty7			
tty8 tty9			
ttyl1 ttyS1			

Configuring serial terminal in RHEL 7

To configure serial terminal in RHEL 7:

1. Add, or update the following lines to /etc/default/grub:

GRUB CMDLINE LINUX DEFAULT="console=tty0 console=ttyS0,115200n8"

GRUB_TERMINAL="console serial"

GRUB_SERIAL_COMMAND="serial --speed=115200 --unit=0 --word=8 --parity=no --stop=1"

GRUB_CMDLINE_LINUX_DEFAULT applies this configuration only to the default menu entry, use GRUB_CMDLINE_LINUX to apply it to all the menu entries.

Each line should only appear once within the /etc/default/grub. If the line already exists, then modify it to avoid another copy. Therefore, only one GRUB CMDLINE LINUX DEFAULT line is allowed.

- 2. Rebuild the /boot/grub.cfg configuration file by running the grub2-mkconfig -o command as follows:
 - on BIOS-based systems:

```
~]# grub2-mkconfig -o /boot/grub2/grub.cfg
```

• on UEFI-based systems:

~] # grub2-mkconfig -o /boot/efi/EFI/redhat/grub.cfg

For more information, see the RHEL 7 System Administrator's Guide at redhat.com.

Controlling GRUB from serial console

You can configure GRUB to use the serial console instead of the VGA console. This allows you to interrupt the boot process and choose a different kernel or add kernel parameters, for example, to boot into single user mode.



To configure GRUB to use serial console, comment out the splash image and add the serial and terminal options to grub.conf:

[root@localhost ~]# cat /boot/grub.conf
grub.conf generated by anaconda
#
Note that you do not have to rerun grub after making changes to this file
NOTICE: You have a /boot partition. This means that
all kernel and initrd paths are relative to /boot/, eg.
root (hd0,0)
<pre># kernel /vmlinuz-version ro root=/dev/hda2</pre>
initrd /initrd-version.img
#boot=/dev/hda
default=0
timeout=10
#splashimage=(hd0,0)/grub/splash.xpm.gz
serialunit=0speed=1152001

NOTE: Restart the system for the settings to take effect.

Supported SSH cryptography schemes

To communicate with iDRAC using SSH protocol, it supports multiple cryptography schemes listed in the following table.

Table 19. SSH cryptography schemes

Scheme Type	Algorithms
Asymmetric Cryptography	
Public key	ssh-rsa ecdsa-sha2-nistp256
Symmetric Cryptography	
Key Exchange	curve25519-sha256@libssh.org
	ecdh-sha2-nistp256
	ecdh-sha2-nistp384
	ecdh-sha2-nistp521
	diffie-hellman-group-exchange-sha256

Setting up iDRAC communication 139



Table 19. SSH cryptography schemes (continued)

Scheme Type	Algorithms
	diffie-hellman-group14-sha1
Encryption	chacha20-poly1305@openssh.com aes128-ctr aes192-ctr aes256-ctr aes128-gcm@openssh.com aes256-gcm@openssh.com
MAC	hmac-sha1 hmac-ripemd160 umac-64@openssh.com
Compression	INOTIE

NOTE: If you enable OpenSSH 7.0 or later, DSA public key support is disabled. To ensure better security for iDRAC, Dell recommends not enabling DSA public key support.

Using public key authentication for SSH

iDRAC supports the Public Key Authentication (PKA) over SSH. This is a licensed feature. When the PKA over SSH is set up and used correctly, you must enter the user name while logging into iDRAC. This is useful for setting up automated scripts that perform various functions. The uploaded keys must be in RFC 4716 or OpenSSH format. Else, you must convert the keys into that format.

In any scenario, a pair of private and public key must be generated on the management station. The public key is uploaded to iDRAC local user and private key is used by the SSH client to establish the trust relationship between the management station and iDRAC.

You can generate the public or private key pair using:

- PuTTY Key Generator application for clients running Windows
- ssh-keygen CLI for clients running Linux.

CAUTION: This privilege is normally reserved for users who are members of the Administrator user group on iDRAC. However, users in the 'Custom' user group can be assigned this privilege. A user with this privilege can modify any user's configuration. This includes creation or deletion of any user, SSH Key management for users, and so on. For these reasons, assign this privilege carefully.

CAUTION: The capability to upload, view, and/ or delete SSH keys is based on the 'Configure Users' user privilege. This privilege allows user(s) to configure another user's SSH key. You should grant this privilege carefully.

Generating public keys for Windows

To use the PuTTY Key Generator application to create the basic key:

- 1. Start the application and select RSA for the key type.
- 2. Enter the number of bits for the key. The number of bits must be between 2048 and 4096 bits.
- Click Generate and move the mouse in the window as directed. The keys are generated.
- 4. You can modify the key comment field.
- 5. Enter a passphrase to secure the key.
- 6. Save the public and private key.



Generating public keys for Linux

To use the *ssh-keygen* application to create the basic key, open a terminal window and at the shell prompt, enter *ssh-keygen* -t *rsa* -b 2048 -C testing

where:

- -t is rsa.
- -b specifies the bit encryption size between 2048 and 4096.
- -C allows modifying the public key comment and is optional.

(i) NOTE: The options are case-sensitive.

Follow the instructions. After the command executes, upload the public file.

CAUTION: Keys generated from the Linux management station using ssh-keygen are in non-4716 format. Convert the keys into the 4716 format using ssh-keygen -e -f /root/.ssh/id_rsa.pub > std_rsa.pub. Do not change the permissions of the key file. The conversion must be done using default permissions.

(i) NOTE: iDRAC does not support ssh-agent forward of keys.

Uploading SSH keys

You can upload up to four public keys *per user* to use over an SSH interface. Before adding the public keys, make sure that you view the keys if they are set up, so that a key is not accidentally overwritten.

When adding new public keys, make sure that the existing keys are not at the index where the new key is added. iDRAC does not perform checks to make sure previous key(s) are deleted before a new key(s) are added. When a new key is added, it is usable if the SSH interface is enabled.

Uploading SSH keys using web interface

To upload the SSH keys:

- In the iDRAC Web interface, go to iDRAC Settings > Users > Local Users. The Local Users page is displayed.
- 2. In the User ID column, click a user ID number. The Users Main Menu page is displayed.
- **3.** Under SSH Key Configurations, select Upload SSH Key(s) and click Next. The Upload SSH Key(s) page is displayed.
- 4. Upload the SSH keys in one of the following ways:
 - Upload the key file.
 - Copy the contents of the key file into the text box

For more information, see iDRAC Online Help.

5. Click Apply.

Uploading SSH keys using RACADM

To upload the SSH keys, run the following command:

(i) NOTE: You cannot upload and copy a key at the same time.

- For local RACADM: racadm sshpkauth -i <2 to 16> -k <1 to 4> -f <filename>
- From remote RACADM using or SSH: racadm sshpkauth -i <2 to 16> -k <1 to 4> -t <key-text>

For example, to upload a valid key to iDRAC User ID 2 in the first key space using a file, run the following command:

\$ racadm sshpkauth -i 2 -k 1 -f pkkey.key

(i) **NOTE:** The -f option is not supported on ssh/serial RACADM.



Viewing SSH keys

You can view the keys that are uploaded to iDRAC.

Viewing SSH keys using web interface

To view the SSH keys:

- In Web interface, go to iDRAC Settings > Users. The Local Users page is displayed.
- 2. In the User ID column, click a user ID number. The Users Main Menu page is displayed.
- **3.** Under SSH Key Configurations, select View/Remove SSH Key(s) and click Next. The View/Remove SSH Key(s) page is displayed with the key details.

Deleting SSH keys

Before deleting the public keys, make sure that you view the keys if they are set up, so that a key is not accidentally deleted.

Deleting SSH keys using web interface

To delete the SSH key(s):

- 1. In Web interface, go to **iDRAC Settings** > **Users**. The **Local Users** page is displayed.
- In the ID column, select a user ID number, click Edit. The Edit User page is displayed.
- **3.** Under **SSH Key Configurations**, select a SSH Key and click **Edit**. The **SSH Key** page displays the **Edit From** details.
- Select Remove for the key(s) you want to delete, and click Apply. The selected key(s) is deleted.

Deleting SSH keys using RACADM

To delete the SSH key(s), run the following commands:

- Specific key racadm sshpkauth -i <2 to 16> -d -k <1 to 4>
- All keys racadm sshpkauth -i <2 to 16> -d -k all



Configuring user accounts and privileges

You can setup user accounts with specific privileges (*role-based authority*) to manage your system using iDRAC and maintain system security. By default iDRAC is configured with a local administrator account. The default iDRAC user name and password are provided with the system badge. As an administrator, you can setup user accounts to allow other users to access iDRAC. For more information see the documentation for the server.

You can setup local users or use directory services such as Microsoft Active Directory or LDAP to setup user accounts. Using a directory service provides a central location for managing authorized user accounts.

iDRAC supports role-based access to users with a set of associated privileges. The roles are administrator, operator, read only, or none. The role defines the maximum privileges available.

Topics:

- iDRAC user roles and privileges
- Recommended characters in user names and passwords
- Configuring local users
- Configuring Active Directory users
- Configuring generic LDAP users

iDRAC user roles and privileges

The iDRAC role and privilege names have changed from earlier generation of servers. The role names are:

Table 20. iDRAC roles

Current Generation	Prior Generation	Privileges
Administrator	Administrator	Login, Configure, Configure Users, Logs, System Control, Access Virtual Console, Access Virtual Media, System Operations, Debug
Operator	Power User	Login, Configure, System Control, Access Virtual Console, Access Virtual Media, System Operations, Debug
Read Only	Guest User	Login
None	None	None

The following table describes the user privileges:

Table 21. iDRAC user privileges

Current Generation	Prior Generation	Description
Login	Login to iDRAC	Enables the user to log in to iDRAC.
Configure	Configure iDRAC	Enables the user to configure iDRAC. With this privilege, a user can also configure power management, virtual console, virtual media, licenses, system settings, storage devices, BIOS settings, SCP and so on.
i NOTE: The adminis	strator role overrides all the priv	vileges from the other components such as BIOS setup password.
Configure Users	Configure Users	Enables the user to allow specific users to access the system.
Logs	Clear Logs	Enables the user to clear only the System Event Log (SEL).

Configuring user accounts and privileges 143



Table 21. iDRAC user privileges (continued)

Current Generation	Prior Generation	Description
System Control	Control and configure system	Allows power cycling the host system.
Access Virtual Console	Access Virtual Console Redirection (for blade servers) Access Virtual Console (for rack and tower servers)	Enables the user to run Virtual Console.
Access Virtual Media	Access Virtual Media	Enables the user to run and use Virtual Media.
System Operations	Test Alerts	Allows user initiated and generated events, and information is sent as an asynchronous notification and logged.
Debug	Execute Diagnostic Commands	Enables the user to run diagnostic commands.

Recommended characters in user names and passwords

This section provides details about the recommended characters while creating and using user names and passwords.

(i) NOTE: The password must include one uppercase and one lower case letter, one number and a special character.

Use the following characters while creating user names and passwords:

Table 22. Recommended characters for user names

Characters	Length
0-9	1–16
A-Z	
a-z	
- ! # \$ % & () * ; ? [\] ^ _ ` { } ~+ <= >	

Table 23. Recommended characters for passwords

Characters	Length
0-9	1-40
A-Z	
a-z	
'-!"#\$%&()*,./:;?@[\]^_`{ }~+<=>	

NOTE: You may be able to create user names and passwords that include other characters. However, to ensure compatibility with all interfaces, Dell recommends using only the characters listed here.

NOTE: The characters allowed in user names and passwords for network shares are determined by the network-share type. iDRAC supports valid characters for network share credentials as defined by the share type, except <, >, and , (comma).

() NOTE: To improve security, it is recommended to use complex passwords that have eight or more characters and include lowercase alphabets, uppercase alphabets, numbers, and special characters. It is also recommended to regularly change the passwords, if possible.

144 Configuring user accounts and privileges



Configuring local users

You can configure up to 16 local users in iDRAC with specific access permissions. Before you create an iDRAC user, verify if any current users exist. You can set user names, passwords, and roles with the privileges for these users. The user names and passwords can be changed using any of the iDRAC secured interfaces (that is, web interface, RACADM or WSMan). You can also enable or disable SNMPv3 authentication for each user.

Configuring local users using iDRAC web interface

To add and configure local iDRAC users:

(i) NOTE: You must have Configure Users permission to create an iDRAC user.

- In the iDRAC Web interface, go to iDRAC Settings > User. The Local Users page is displayed.
- 2. In the User ID column, select a user ID number and click Edit.

(i) NOTE: User 1 is reserved for the IPMI anonymous user and you cannot change this configuration.

The User Configuration page is displayed.

3. Add User Account Settings and Advanced Settings details to configure the user account.

() **NOTE:** Enable the user ID and specify the user name, password, and user role (access privileges) for the user. You can also enable LAN privilege level, Serial port privilege level, serial over LAN status, SNMPv3 authentication, authentication type and the privacy type for the user. For more information about the options, see the *iDRAC Online Help*.

4. Click **Save**. The user is created with the required privileges.

Configuring local users using RACADM

(i) NOTE: You must be logged in as user **root** to execute RACADM commands on a remote Linux system.

You can configure single or multiple iDRAC users using RACADM.

To configure multiple iDRAC users with identical configuration settings, follow these procedures:

- Use the RACADM examples in this section as a guide to create a batch file of RACADM commands and then execute the batch file on each managed system.
- Create the iDRAC configuration file and execute the racadm set command on each managed system using the same configuration file.

If you are configuring a new iDRAC or if you have used the racadm racresetcfg command, then check for the default iDRAC user name and password on the system badge. The racadm racresetcfg command resets the iDRAC to the default values.

NOTE: If SEKM is enabled on the server, then disable SEKM using the racadm sekm disable command before using this command. This can avoid any storage devices being locked out which are secured by iDRAC, if SEKM settings are erased from iDRAC by executing this command.

(i) NOTE: Users can be enabled and disabled over time. As a result, a user may have a different index number on each iDRAC.

To verify if a user exists, type the following command once for each index (1–16):

racadm get iDRAC.Users.<index>.UserName

Several parameters and object IDs are displayed with their current values. The key field is iDRAC.Users.UserName=. If a user name is displayed after =, that index number is taken.

(i) NOTE: You can utilize

racadm get -f <myfile.cfg>

Fis. <u>532</u> Mov. <u>32</u>

and view or edit the

myfile.cfg

file, which includes all iDRAC configuration parameters.

To enable SNMP v3 authentication for a user, use **SNMPv3AuthenticationType**, **SNMPv3Enable**, **SNMPv3PrivacyType** objects. For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

If you use the Server Configuration Profile file to configure users, use the **AuthenticationProtocol**, **ProtocolEnable**, and **PrivacyProtocol** attributes to enable SNMPv3 authentication.

Adding iDRAC user using RACADM

1. Set the index and user name.

racadm set idrac.users.<index>.username <user name>

Parameter	Description	
<index></index>	Unique index of the user	
<user_name></user_name>	User name	

2. Set the password.

racadm set idrac.users.<index>.password <password>

3. Set the user privileges.

For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

4. Enable the user.

racadm set.idrac.users.<index>.enable 1

To verify, use the following command:

racadm get idrac.users.<index>

For more information, see the iDRAC RACADM CLI Guide available at https://www.dell.com/idracmanuals.

Enabling iDRAC user with permissions

To enable a user with specific administrative permissions (role-based authority):

1. Locate an available user index.

racadm get iDRAC.Users <index>

2. Type the following commands with the new user name and password.

racadm set iDRAC.Users.<index>.Privilege <user privilege bit mask value>

NOTE: The default privilege value is 0, which indicates the user has no privileges enabled. For a list of valid bit-mask values for specific user privileges, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.



Configuring Active Directory users

If your company uses the Microsoft Active Directory software, you can configure the software to provide access to iDRAC, allowing you to add and control iDRAC user privileges to your existing users in your directory service. This is a licensed feature.

You can configure user authentication through Active Directory to log in to the iDRAC. You can also provide role-based authority, which enables an administrator to configure specific privileges for each user.

NOTE: For any deployment done via MX Template and CA validation is enabled within template, the user must upload CA certificates at first login or before changing Authentication Service from LDAP to Active Directory or vice versa.

Prerequisites for using Active Directory authentication for iDRAC

To use the Active Directory authentication feature of iDRAC, make sure that you have:

- Deployed an Active Directory infrastructure. See the Microsoft website for more information.
- Integrated PKI into the Active Directory infrastructure. iDRAC uses the standard Public Key Infrastructure (PKI) mechanism to authenticate securely into the Active Directory. See the Microsoft website for more information.
- Enabled the Secure Socket Layer (SSL) on all domain controllers that iDRAC connects to for authenticating to all the domain controllers.

Enabling SSL on domain controller

When iDRAC authenticates users with an Active Directory domain controller, it starts an SSL session with the domain controller. At this time, the domain controller must publish a certificate signed by the Certificate Authority (CA)—the root certificate of which is also uploaded into iDRAC. For iDRAC to authenticate to *any* domain controller—whether it is the root or the child domain controller—that domain controller must have an SSL-enabled certificate signed by the domain's CA.

If you are using Microsoft Enterprise Root CA to *automatically* assign all your domain controllers to an SSL certificate, you must:

- 1. Install the SSL certificate on each domain controller.
- 2. Export the Domain Controller Root CA Certificate to iDRAC.
- 3. Import iDRAC Firmware SSL Certificate.

Installing SSL certificate for each domain controller

To install the SSL certificate for each controller:

- 1. Click Start > Administrative Tools > Domain Security Policy.
- 2. Expand the Public Key Policies folder, right-click Automatic Certificate Request Settings and click Automatic Certificate Request.
 - The Automatic Certificate Request Setup Wizard is displayed.
- 3. Click Next and select Domain Controller.
- 4. Click Next and click Finish. The SSL certificate is installed.

Exporting domain controller root CA certificate to iDRAC

To export the domain controller root CA certificate to iDRAC:

- 1. Locate the domain controller that is running the Microsoft Enterprise CA service.
- 2. Click Start > Run.
- 3. Enter mmc and click OK.
- 4. In the Console 1 (MMC) window, click File (or Console) and select Add/Remove Snap-in.
- 5. In the Add/Remove Snap-In window, click Add.
- 6. In the Standalone Snap-In window, select Certificates and click Add.
- 7. Select Computer and click Next.
- 8. Select Local Computer, click Finish, and click OK.
- 9. In the Console 1 window, go to Certificates Personal Certificates folder.



- 10. Locate and right-click the root CA certificate, select All Tasks, and click Export....
- 11. In the Certificate Export Wizard, click Next, and select No do not export the private key.
- 12. Click Next and select Base-64 encoded X.509 (.cer) as the format.
- 13. Click **Next** and save the certificate to a directory on your system.
- 14. Upload the certificate you saved in step 13 to iDRAC.

Importing iDRAC firmware SSL certificate

iDRAC SSL certificate is the identical certificate used for iDRAC Web server. All iDRAC controllers are shipped with a default self-signed certificate.

If the Active Directory Server is set to authenticate the client during an SSL session initialization phase, you need to upload iDRAC Server certificate to the Active Directory Domain controller. This additional step is not required if the Active Directory does not perform a client authentication during an SSL session's initialization phase.

NOTE: If iDRAC firmware SSL certificate is CA-signed and the certificate of that CA is already in the domain controller's Trusted Root Certificate Authority list, do not perform the steps in this section.

To import iDRAC firmware SSL certificate to all domain controller trusted certificate lists:

1. Download iDRAC SSL certificate using the following RACADM command:

racadm sslcertdownload -t 1 -f <RAC SSL certificate>

- On the domain controller, open an MMC Console window and select Certificates > Trusted Root Certification Authorities.
- 3. Right-click Certificates, select All Tasks and click Import.
- 4. Click Next and browse to the SSL certificate file.
- 5. Install iDRAC SSL Certificate in each domain controller's Trusted Root Certification Authority.

If you have installed your own certificate, make sure that the CA signing your certificate is in the **Trusted Root Certification Authority** list. If the Authority is not in the list, you must install it on all your domain controllers.

- 6. Click **Next** and select whether you want Windows to automatically select the certificate store based on the type of certificate, or browse to a store of your choice.
- 7. Click Finish and click OK. The iDRAC firmware SSL certificate is imported to all domain controller trusted certificate lists.

Supported Active Directory authentication mechanisms

You can use Active Directory to define iDRAC user access using two methods:

- Standard schema solution, which uses Microsoft's default Active Directory group objects only.
- *Extended schema* solution, which has customized Active Directory objects. All the access control objects are maintained in Active Directory. It provides maximum flexibility to configure user access on different iDRACs with varying privilege levels.

Standard schema Active Directory overview

As shown in the following figure, using standard schema for Active Directory integration requires configuration on both Active Directory and iDRAC.





Figure 1. Configuration of iDRAC with active directory standard schema

In Active Directory, a standard group object is used as a role group. A user who has iDRAC access is a member of the role group. To give this user access to a specific iDRAC, the role group name and its domain name need to be configured on the specific iDRAC. The role and the privilege level are defined on each iDRAC and not in the Active Directory. You can configure up to 15 role groups in each iDRAC. Table reference no shows the default role group privileges.

Role Groups	Default Privilege Level	Permissions Granted	Bit Mask
Role Group 1	None	Log in to iDRAC, Configure iDRAC, Configure Users, Clear Logs, Execute Server Control Commands, Access Virtual Console, Access Virtual Media, Test Alerts, Execute Diagnostic Commands	0x000001ff
Role Group 2	None	Log in to iDRAC, Configure iDRAC, Execute Server Control Commands, Access Virtual Console, Access Virtual Media, Test Alerts, Execute Diagnostic Commands	0x00000f9
Role Group 3	None	Log in to iDRAC	0x0000001
Role Group 4	None	No assigned permissions	0×0000000
Role Group 5	None	No assigned permissions	0×0000000
			1

Table 24. Default role group privileges

(i) NOTE: The Bit Mask values are used only when setting Standard Schema with the RACADM.

Single domain versus multiple domain scenarios

If all the login users and role groups, including the nested groups, are in the same domain, then only the domain controllers' addresses must be configured on iDRAC. In this single domain scenario, any group type is supported.

If all the login users and role groups, or any of the nested groups, are from multiple domains, then Global Catalog server addresses must be configured on iDRAC. In this multiple domain scenario, all the role groups and nested groups, if any, must be a Universal Group type.



Configuring Standard schema Active Directory

Before configuring the standard schema Active Directory, ensure that:

- You have the iDRAC Enterprise or Datacenter license.
- The configuration is performed on a server that is used as the Domain Controller.
- The dat, time and time zone on the server are correct.
- The iDRAC network settings are configured, or in iDRAC web interface go to iDRAC Settings > Connectivity > Network > Common Settings to configure the network settings.

To configure iDRAC for an Active Directory login access:

- 1. On an Active Directory server (domain controller), open the Active Directory Users and Computers Snap-in.
- 2. Create the iDRAC groups and users.
- 3. Configure the group name, domain name, and the role privileges on iDRAC using the iDRAC web interface or RACADM.

Configuring Active Directory with Standard schema using iDRAC web interface

(i) NOTE: For information about the various fields, see the *iDRAC Online Help*.

- In the iDRAC web interface, go to iDRAC Settings > Users > Directory Services. The Directory Service page is displayed.
- Select the Microsoft Active Directory option and then click Edit. The Active Directory Configuration and Management page is displayed.
- 3. Click Configure Active Directory. The Active Directory Configuration and Management Step 1 of 4 page is displayed.
- 4. Optionally, enable certificate validation and upload the CA-signed digital certificate used during initiation of SSL connections when communicating with the Active Directory (AD) server. For this, the Domain Controllers and Global Catalog FQDN must be specified. This is done in the next steps. And hence the DNS should be configured properly in the network settings.
- 5. Click Next.

The Active Directory Configuration and Management Step 2 of 4 page is displayed.

6. Enable Active Directory and specify the location information about Active Directory servers and user accounts. Also, specify the time iDRAC must wait for responses from Active Directory during iDRAC login.

NOTE: If certificate validation is enabled, specify the Domain Controller Server addresses and the Global Catalog FQDN. Make sure that DNS is configured correctly under **iDRAC Settings** > **Network**.

- 7. Click Next. The Active Directory Configuration and Management Step 3 of 4 page is displayed.
- Select Standard Schema and click Next.
 The Active Directory Configuration and Management Step 4a of 4 page is displayed.
- 9. Enter the location of Active Directory global catalog server(s) and specify privilege groups used to authorize users.
- Click a Role Group to configure the control authorization policy for users under the standard schema mode. The Active Directory Configuration and Management Step 4b of 4 page is displayed.
- 11. Specify the privileges and click Apply.

The settings are applied and the Active Directory Configuration and Management Step 4a of 4 page is displayed.

12. Click **Finish**. The Active Directory settings for standard schema are configured.

Configuring Active Directory with Standard schema using RACADM

1. Use the following commands:

racadm	set	iDRAC.ActiveDirectory.Enable 1
racadm	set	iDRAC.ActiveDirectory.Schema 2
racadm	set	iDRAC.ADGroup.Name <common group="" name="" of="" role="" the=""></common>
racadm	set	iDRAC.ADGroup.Domain <fully domain="" name="" qualified=""></fully>
racadm	set	iDRAC.ADGroup.Privilege <bit-mask for="" permissions="" rolegroup="" specific="" value=""></bit-mask>
racadm	set	iDRAC.ActiveDirectory.DomainController1 <fully domain="" ip<="" name="" or="" qualified="" td=""></fully>

150 Configuring user accounts and privileges



```
address of the domain controller>
racadm set iDRAC.ActiveDirectory.DomainController2 <fully qualified domain name or IP
address of the domain controller>
racadm set iDRAC.ActiveDirectory.DomainController3 <fully qualified domain name or IP
address of the domain controller>
racadm set iDRAC.ActiveDirectory.GlobalCatalog1 <fully qualified domain name or IP
address of the domain controller>
racadm set iDRAC.ActiveDirectory.GlobalCatalog2 <fully qualified domain name or IP
address of the domain controller>
racadm set iDRAC.ActiveDirectory.GlobalCatalog2 <fully qualified domain name or IP
address of the domain controller>
racadm set iDRAC.ActiveDirectory.GlobalCatalog3 <fully qualified domain name or IP
address of the domain controller>
```

- Enter the Fully Qualified Domain Name (FQDN) of the domain controller, not the FQDN of the domain. For example, enter servername.dell.com instead of dell.com.
- For bit-mask values for specific Role Group permissions, see Default role group privileges.
- You must provide at least one of the three domain controller addresses. iDRAC attempts to connect to each of the configured addresses one-by-one until it makes a successful connection. With Standard Schema, these are the addresses of the domain controllers where the user accounts and the role groups are located.
- The Global Catalog server is only required for standard schema when the user accounts and role groups are in different domains. In multiple domain case, only the Universal Group can be used.
- If certificate validation is enabled, the FQDN or IP address that you specify in this field must match the Subject or Subject Alternative Name field of your domain controller certificate.
- To disable the certificate validation during SSL handshake, use the following command:

racadm set iDRAC.ActiveDirectory.CertValidationEnable 0

In this case, no Certificate Authority (CA) certificate needs to be uploaded.

To enforce the certificate validation during SSL handshake (optional), use the following command:

racadm set iDRAC.ActiveDirectory.CertValidationEnable 1

In this case, you must upload the CA certificate using the following command:

racadm sslcertupload -t 0x2 -f <ADS root CA certificate>

NOTE: If certificate validation is enabled, specify the Domain Controller Server addresses and the Global Catalog FQDN. Ensure that DNS is configured correctly under **Overview** > **iDRAC Settings** > **Network**.

Using the following RACADM command may be optional.

racadm sslcertdownload -t 1 -f <RAC SSL certificate>

2. If DHCP is enabled on iDRAC and you want to use the DNS provided by the DHCP server, enter the following command:

racadm set iDRAC.IPv4.DNSFromDHCP 1

3. If DHCP is disabled on iDRAC or you want manually enter the DNS IP address, enter the following RACADM command:

```
racadm set iDRAC.IPv4.DNSFromDHCP 0
racadm set iDRAC.IPv4.DNSFromDHCP.DNS1 <primary DNS IP address>
racadm set iDRAC.IPv4.DNSFromDHCP.DNS2 <secondary DNS IP address>
```

4. If you want to configure a list of user domains so that you only need to enter the user name when logging in to the web interface, use the following command:

```
racadm set iDRAC.UserDomain.<index>.Name <fully qualified domain name or IP Address
of the domain controller>
```

You can configure up to 40 user domains with index numbers between 1 and 40.

Extended schema Active Directory overview

Using the extended schema solution requires the Active Directory schema extension.

Configuring user accounts and privileges 151



Best practices for extended schema

The extended schema uses Dell association objects to join iDRAC and permission. This allows you to use iDRAC based on the overall permissions granted. The default Access Control List (ACL) of Dell Association objects allows Self and Domain Administrators to manage the permissions and scope of iDRAC objects.

By default, the Dell Association objects do not inherit all permissions from the parent Active Directory objects. If you enable inheritance for the Dell Association object, the inherited permissions for that association object are granted to the selected users and groups. This may result in unintended privileges being provided to the iDRAC.

To use the Extended Schema securely, Dell recommends not enabling inheritance on Dell Association objects within the extended schema implementation.

Active directory schema extensions

The Active Directory data is a distributed database of *attributes* and *classes*. The Active Directory schema includes the rules that determine the type of data that can be added or included in the database. The user class is one example of a *class* that is stored in the database. Some example user class attributes can include the user's first name, last name, phone number, and so on. You can extend the Active Directory database by adding your own unique *attributes* and *classes* for specific requirements. Dell has extended the schema to include the necessary changes to support remote management authentication and authorization using Active Directory.

Each *attribute* or *class* that is added to an existing Active Directory Schema must be defined with a unique ID. To maintain unique IDs across the industry, Microsoft maintains a database of Active Directory Object Identifiers (OIDs) so that when companies add extensions to the schema, they can be guaranteed to be unique and not to conflict with each other. To extend the schema in Microsoft's Active Directory, Dell received unique OIDs, unique name extensions, and uniquely linked attribute IDs for the attributes and classes that are added into the directory service:

- Extension is: dell
- Base OID is: 1.2.840.113556.1.8000.1280
- RAC LinkID range is: 12070 to 12079

Overview of iDRAC schema extensions

Dell has extended the schema to include an *Association, Device,* and *Privilege* property. The *Association* property is used to link together the users or groups with a specific set of privileges to one or more iDRAC devices. This model provides an administrator maximum flexibility over the different combinations of users, iDRAC privileges, and iDRAC devices on the network without much complexity.

For each physical iDRAC device on the network that you want to integrate with Active Directory for authentication and authorization, create at least one association object and one iDRAC device object. You can create multiple association objects, and each association object can be linked to as many users, groups of users, or iDRAC device objects as required. The users and iDRAC user groups can be members of any domain in the enterprise.

However, each association object can be linked (or, may link users, groups of users, or iDRAC device objects) to only one privilege object. This example allows an administrator to control each user's privileges on specific iDRAC devices.

iDRAC device object is the link to iDRAC firmware for querying Active Directory for authentication and authorization. When iDRAC is added to the network, the administrator must configure iDRAC and its device object with its Active Directory name so that users can perform authentication and authorization with Active Directory. Additionally, the administrator must add iDRAC to at least one association object for users to authenticate.

The following figure shows that the association object provides the connection that is needed for the authentication and authorization.





Figure 2. Typical setup for active directory objects

You can create as many or as few association objects as required. However, you must create at least one Association Object, and you must have one iDRAC Device Object for each iDRAC device on the network that you want to integrate with Active Directory for Authentication and Authorization with iDRAC.

The Association Object allows for as many or as few users and/or groups as well as iDRAC Device Objects. However, the Association Object only includes one Privilege Object per Association Object. The Association Object connects the Users who have Privileges on iDRAC devices.

The Dell extension to the ADUC MMC Snap-in only allows associating the Privilege Object and iDRAC Objects from the same domain with the Association Object. The Dell extension does not allow a group or an iDRAC object from other domains to be added as a product member of the Association Object.

When adding Universal Groups from separate domains, create an Association Object with Universal Scope. The Default Association objects created by the Dell Schema Extender Utility are Domain Local Groups and they do not work with Universal Groups from other domains.

Users, user groups, or nested user groups from any domain can be added into the Association Object. Extended Schema solutions support any user group type and any user group nesting across multiple domains allowed by Microsoft Active Directory.

Accumulating privileges using Extended Schema

The Extended Schema Authentication mechanism supports Privilege Accumulation from different privilege objects associated with the same user through different Association Objects. In other words, Extended Schema Authentication accumulates privileges to allow the user the super set of all assigned privileges corresponding to the different privilege objects associated with the same user.

The following figure provides an example of accumulating privileges using Extended Schema.



Figure 3. Privilege accumulation for a user



The figure shows two Association Objects—A01 and A02. User1 is associated to iDRAC2 through both association objects.

Extended Schema Authentication accumulates privileges to allow the user the maximum set of privileges possible considering the assigned privileges of the different privilege objects associated to the same user.

In this example, User1 has both Priv1 and Priv2 privileges on iDRAC2. User1 has Priv1 privileges on iDRAC1 only. User2 has Priv1 privileges on both iDRAC1 and iDRAC2. In addition, this figure shows that User1 can be in a different domain and can be a member of a group.

Configuring Extended schema Active Directory

To configure Active Directory to access iDRAC:

- 1. Extend the Active Directory schema.
- 2. Extend the Active Directory Users and Computers Snap-in.
- **3.** Add iDRAC users and their privileges to Active Directory.
- 4. Configure iDRAC Active Directory properties using iDRAC Web interface or RACADM.

Extending Active Directory schema

Extending your Active Directory schema adds a Dell organizational unit, schema classes and attributes, and example privileges and association objects to the Active Directory schema. Before you extend the schema, make sure that you have the Schema Admin privileges on the Schema Master FSMO-Role-Owner of the domain forest.

NOTE: The schema extension for this product is different from the previous generations. The earlier schema does not work with this product.

(i) NOTE: Extending the new schema has no impact on previous versions of the product.

You can extend your schema using one of the following methods:

- Dell Schema Extender utility
- LDIF script file

If you use the LDIF script file, the Dell organizational unit is not added to the schema.

The LDIF files and Dell Schema Extender are on your *Dell Systems Management Tools and Documentation* DVD in the following respective directories:

- DVDdrive:\SYSMGMT\ManagementStation\support\OMActiveDirectory_Tools\Remote_Management_Adv anced\LDIF_Files
- OVDdrive>:

\SYSMGMT\ManagementStation\support\OMActiveDirectory_Tools\Remote_Management_Advanced\Sch ema Extender

To use the LDIF files, see the instructions in the readme included in the LDIF_Files directory.

You can copy and run the Schema Extender or LDIF files from any location.

Using Dell Schema Extender

CAUTION: The Dell Schema Extender uses the SchemaExtenderOem.ini file. To make sure that the Dell Schema Extender utility functions properly, do not modify the name of this file.

- 1. In the Welcome screen, click Next.
- 2. Read and understand the warning and click Next.
- 3. Select Use Current Log In Credentials or enter a user name and password with schema administrator rights.
- 4. Click Next to run the Dell Schema Extender.
- 5. Click Finish.

The schema is extended. To verify the schema extension, use the MMC and the Active Directory Schema Snap-in to verify that the Classes and attributes exist. See the Microsoft documentation for details about using the MMC and the Active Directory Schema Snap-in.



Classes and attributes

Table 25. Class definitions for classes added to the active directory schema

Class Name	Assigned Object Identification Number (OID)
delliDRACDevice	1.2.840.113556.1.8000.1280.1.7.1.1
delliDRACAssociation	1.2.840.113556.1.8000.1280.1.7.1.2
dellRAC4Privileges	1.2.840.113556.1.8000.1280.1.1.1.3
dellPrivileges	1.2.840.113556.1.8000.1280.1.1.1.4
dellProduct	1.2.840.113556.1.8000.1280.1.1.1.5

Table 26. DelliDRACdevice class

OID	1.2.840.113556.1.8000.1280.1.7.1.1
Description	Represents the Dell iDRAC device. iDRAC must be configured as delliDRACDevice in Active Directory. This configuration enables iDRAC to send Lightweight Directory Access Protocol (LDAP) queries to Active Directory.
Class Type	Structural Class
SuperClasses	dellProduct
Attributes	dellSchemaVersion dellRacType

Table 27. delliDRACAssociationObject class

OID	1.2.840.113556.1.8000.1280.1.7.1.2
Description	Represents the Dell Association Object. The Association Object provides the connection between the users and the devices.
Class Type	Structural Class
SuperClasses	Group
Attributes	dellProductMembers dellPrivilegeMember

Table 28. dellRAC4Privileges class

OID	1.2.840.113556.1.8000.1280.1.1.1.3
Description	Defines the privileges (Authorization Rights) for iDRAC
Class Type	Auxiliary Class
SuperClasses	None
Attributes	delllsLoginUser
	dellIsCardConfigAdmin
	dellIsUserConfigAdmin



Table 28. dellRAC4Privileges class (continued)

1.2.840.113556.1.8000.1280.1.1.1.3
dellIsLogClearAdmin
dellIsServerResetUser
dellIsConsoleRedirectUser
dellIsVirtualMediaUser
dellIsTestAlertUser
dellIsDebugCommandAdmin

Table 29. dellPrivileges class

OID	1.2.840.113556.1.8000.1280.1.1.1.4
Description	Used as a container Class for the Dell Privileges (Authorization Rights).
Class Type	Structural Class
SuperClasses	User
Attributes	dellRAC4Privileges

Table 30. dellProduct class

OID	1.2.840.113556.1.8000.1280.1.1.1.5
Description	The main class from which all Dell products are derived.
Class Type	Structural Class
SuperClasses	Computer
Attributes	dellAssociationMembers

Table 31. List of attributes added to the active directory schema

Attribute Name/Description	Assigned OID/Syntax Object Identifier	Single Valued
dellPrivilegeMember	1.2.840.113556.1.8000.1280.1.1.2.1	FALSE
List of dellPrivilege Objects that belong to this Attribute.	Distinguished Name (LDAPTYPE_DN 1.3.6.1.4.1.1466.115.121.1.12)	
dellProductMembers	1.2.840.113556.1.8000.1280.1.1.2.2	FALSE
List of dellRacDevice and DelliDRACDevice Objects that belong to this role. This attribute is the forward link to the dellAssociationMembers backward link. Link ID: 12070	Distinguished Name (LDAPTYPE_DN 1.3.6.1.4.1.1466.115.121.1.12)	
delllsLoginUser	1.2.840.113556.1.8000.1280.1.1.2.3	TRUE
TRUE if the user has Login rights on the device.	Boolean (LDAPTYPE_BOOLEAN 1.3.6.1.4.1.1466.115.121.1.7)	
dellIsCardConfigAdmin	1.2.840.113556.1.8000.1280.1.1.2.4	TRUE

Inserido ao protocolo **22.951.206-4** por: **Pedro Henrique Golin Linhares** em: 31/03/2025 12:49. A autenticidade deste documento pode ser validada no endereço: https://www.eprotocolo.pr.gov.br/spiweb/validarDocumento com o código: 692a5ef82beb173ed33dad4ee6505546.



Table 31. List of attributes added to the active directory schema (continued)

Attribute Name/Description	Assigned OID/Syntax Object Identifier	Single Valued	
TRUE if the user has Card Configuration rights on the device.	Boolean (LDAPTYPE_BOOLEAN 1.3.6.1.4.1.1466.115.121.1.7)		
dellIsUserConfigAdmin	1.2.840.113556.1.8000.1280.1.1.2.5	TRUE	
TRUE if the user has User Configuration rights on the device.	Boolean (LDAPTYPE_BOOLEAN 1.3.6.1.4.1.1466.115.121.1.7)		
dellsLogClearAdmin	1.2.840.113556.1.8000.1280.1.1.2.6	TRUE	
TRUE if the user has Log Clearing rights on the device.	Boolean (LDAPTYPE_BOOLEAN 1.3.6.1.4.1.1466.115.121.1.7)		
dellIsServerResetUser	1.2.840.113556.1.8000.1280.1.1.2.7	TRUE	
TRUE if the user has Server Reset rights on the device.	Boolean (LDAPTYPE_BOOLEAN 1.3.6.1.4.1.1466.115.121.1.7)		
dellIsConsoleRedirectUser	1.2.840.113556.1.8000.1280.1.1.2.8	TRUE	
TRUE if the user has Virtual Console rights on the device.	Boolean (LDAPTYPE_BOOLEAN 1.3.6.1.4.1.1466.115.121.1.7)		
dellIsVirtualMediaUser	1.2.840.113556.1.8000.1280.1.1.2.9	TRUE	
TRUE if the user has Virtual Media rights on the device.	Boolean (LDAPTYPE_BOOLEAN 1.3.6.1.4.1.1466.115.121.1.7)		
dellIsTestAlertUser	1.2.840.113556.1.8000.1280.1.1.2.10	TRUE	
TRUE if the user has Test Alert User rights on the device.	Boolean (LDAPTYPE_BOOLEAN 1.3.6.1.4.1.1466.115.121.1.7)		
dellIsDebugCommandAdmin	1.2.840.113556.1.8000.1280.1.1.2.11 TRUE		
TRUE if the user has Debug Command Admin rights on the device.	Boolean (LDAPTYPE_BOOLEAN 1.3.6.1.4.1.1466.115.121.1.7)		
dellSchemaVersion	1.2.840.113556.1.8000.1280.1.1.2.12	TRUE	
The Current Schema Version is used to update the schema.	Case Ignore String (LDAPTYPE_CASEIGNORESTRING 1.2.840.113556.1.4.905)		
dellRacType	1.2.840.113556.1.8000.1280.1.1.2.13	TRUE	
This attribute is the Current RAC Type for the delliDRACDevice object and the backward link to the dellAssociationObjectMembers forward link.	Case Ignore String (LDAPTYPE_CASEIGNORESTRING 1.2.840.113556.1.4.905)		
dellAssociationMembers	1.2.840.113556.1.8000.1280.1.1.2.14	FALSE	
List of dellAssociationObjectMembers that belong to this Product. This attribute is the backward link to the dellProductMembers linked attribute.	Distinguished Name (LDAPTYPE_DN 1.3.6.1.4.1.1466.115.121.1.12)		
Link ID: 12071			



Installing Dell extension to the Active Directory users and computers snap-ir

When you extend the schema in Active Directory, you must also extend the Active Directory Users and Computers Snap-in so the administrator can manage iDRAC devices, users and user groups, iDRAC associations, and iDRAC privileges.

When you install your systems management software using the *Dell Systems Management Tools and Documentation* DVD, you can extend the Snap-in by selecting the **Active Directory Users and Computers Snap-in** option during the installation procedure. See the Dell OpenManage Software Quick Installation Guide for additional instructions about installing systems management software. For 64-bit Windows Operating Systems, the Snap-in installer is located under:

<DVDdrive>:\SYSMGMT\ManagementStation\support\OMActiveDirectory_SnapIn64

For more information about the Active Directory Users and Computers Snap-in, see Microsoft documentation.

Adding iDRAC users and privileges to Active Directory

Using the Dell-extended Active Directory Users and Computers Snap-in, you can add iDRAC users and privileges by creating device, association, and privilege objects. To add each object, perform the following:

- Create an iDRAC device Object
- Create a Privilege Object
- Create an Association Object
- Add objects to an Association Object

Creating iDRAC device object

To create iDRAC device object:

- 1. In the MMC Console Root window, right-click a container.
- Select New > Dell Remote Management Object Advanced. The New Object window is displayed.
- **3.** Enter a name for the new object. The name must be identical to iDRAC name that you enter while configuring Active Directory properties using iDRAC Web interface.
- 4. Select iDRAC Device Object and click OK.

Creating privilege object

To create a privilege object:

(i) NOTE: You must create a privilege object in the same domain as the related association object.

- 1. In the Console Root (MMC) window, right-click a container.
- Select New > Dell Remote Management Object Advanced. The New Object window is displayed.
- 3. Enter a name for the new object.
- 4. Select Privilege Object and click OK.
- 5. Right-click the privilege object that you created, and select Properties.
- 6. Click the Remote Management Privileges tab and assign the privileges for the user or group.

Creating association object

To create association object: **IDRAC** association object is derived from the group and its scope is set to Domain Local.

- 1. In the Console Root (MMC) window, right-click a container.
- Select New > Dell Remote Management Object Advanced. This New Object window is displayed.
- 3. Enter a name for the new object and select Association Object.
- 4. Select the scope for the Association Object and click OK.
- 5. Provide access privileges to the authenticated users for accessing the created association objects.

158 Configuring user accounts and privileges



Providing user access privileges for association objects

To provide access privileges to the authenticated users for accessing the created association objects:

- 1. Go to Administrative Tools > ADSI Edit. The ADSI Edit window is displayed.
- 2. In the right-pane, navigate to the created association object, right-click and select **Properties**.
- 3. In the Security tab, click Add.
- 4. Type Authenticated Users, click Check Names, and click OK. The authenticated users is added to the list of Groups and user names.
- 5. Click OK.

Adding objects to association object

Using the **Association Object Properties** window, you can associate users or user groups, privilege objects, and iDRAC devices or iDRAC device groups.

You can add groups of users and iDRAC devices.

Adding users or user groups

To add users or user groups:

- 1. Right-click the Association Object and select Properties.
- 2. Select the Users tab and click Add.
- 3. Enter the user or user group name and click OK.

Adding privileges

To add privileges:

Click the **Privilege Object** tab to add the privilege object to the association that defines the user's or user group's privileges when authenticating to an iDRAC device. Only one privilege object can be added to an Association Object.

- 1. Select the Privileges Object tab and click Add.
- 2. Enter the privilege object name and click OK.
- **3.** Click the **Privilege Object** tab to add the privilege object to the association that defines the user's or user group's privileges when authenticating to an iDRAC device. Only one privilege object can be added to an Association Object.

Adding iDRAC devices or iDRAC device groups

To add iDRAC devices or iDRAC device groups:

- 1. Select the **Products** tab and click **Add**.
- 2. Enter iDRAC devices or iDRAC device group name and click OK.
- 3. In the Properties window, click Apply and click OK.
- 4. Click the **Products** tab to add one iDRAC device connected to the network that is available for the defined users or user groups. You can add multiple iDRAC devices to an Association Object.

Configuring Active Directory with Extended schema using iDRAC web interface

To configure Active Directory with extended schema using Web interface:

(i) NOTE: For information about the various fields, see the *iDRAC Online Help*.

- 1. In the iDRAC Web interface, go to iDRAC Settings > Users > Directory Services > Microsoft Active Directory. Click Edit
 - The Active Directory Configuration and Management Step 1 of 4 page is displayed.
- 2. Optionally, enable certificate validation and upload the CA-signed digital certificate used during initiation of SSL connections when communicating with the Active Directory (AD) server.



3. Click Next.

- The Active Directory Configuration and Management Step 2 of 4 page is displayed.
- Specify the location information about Active Directory (AD) servers and user accounts. Also, specify the time iDRAC must
 wait for responses from AD during login process.

```
(i) NOTE:
```

- If certificate validation is enabled, specify the Domain Controller Server addresses and the FQDN. Make sure that DNS is configured correctly under iDRAC Settings > Network
- If the user and iDRAC objects are in different domains, then do not select the **User Domain from Login** option. Instead select **Specify a Domain** option and enter the domain name where the iDRAC object is available.
- 5. Click Next. The Active Directory Configuration and Management Step 3 of 4 page is displayed.
- 6. Select Extended Schema and click Next. The Active Directory Configuration and Management Step 4 of 4 page is displayed.
- 7. Enter the name and location of the iDRAC device object in Active Directory (AD) and click **Finish**. The Active Directory settings for extended schema mode is configured.

Configuring Active Directory with Extended schema using RACADM

To configure Active Directory with Extended Schema using the RACADM:

1. Use the following commands:

```
racadm set iDRAC.ActiveDirectory.Enable 1
racadm set iDRAC.ActiveDirectory.Schema 2
racadm set iDRAC.ActiveDirectory.RacName <RAC common name>
racadm set iDRAC.ActiveDirectory.RacDomain <fully qualified rac domain name>
racadm set iDRAC.ActiveDirectory.DomainController1 <fully qualified domain name or IP
address of the domain controller>
racadm set iDRAC.ActiveDirectory.DomainController2 <fully qualified domain name or IP
address of the domain controller>
racadm set iDRAC.ActiveDirectory.DomainController3 <fully qualified domain name or IP
address of the domain controller>
```

- Enter the Fully Qualified Domain Name (FQDN) of the domain controller, not the FQDN of the domain. For example, enter servername.dell.com instead of dell.com.
- You must provide at least one of the three addresses. iDRAC attempts to connect to each of the configured addresses one-by-one until it makes a successful connection. With Extended Schema, these are the FQDN or IP addresses of the domain controllers where this iDRAC device is located.
- To disable the certificate validation during SSL handshake, use the following command:

```
racadm set iDRAC.ActiveDirectory.CertValidationEnable 0
```

In this case, you do not have to upload a CA certificate.

To enforce the certificate validation during SSL handshake (optional):

racadm set iDRAC.ActiveDirectory.CertValidationEnable 1

In this case, you must upload a CA certificate using the following command:

racadm sslcertupload -t 0x2 -f <ADS root CA certificate>

NOTE: If certificate validation is enabled, specify the Domain Controller Server addresses and the FQDN. Ensure that DNS is configured correctly under **iDRAC Settings** > **Network**.

Using the following RACADM command may be optional:

racadm sslcertdownload -t 1 -f <RAC SSL certificate>

2. If DHCP is enabled on iDRAC and you want to use the DNS provided by the DHCP server, enter the following command:

racadm set iDRAC.IPv4.DNSFromDHCP 1



3. If DHCP is disabled in iDRAC or you want to manually input your DNS IP address, enter the following command:

```
racadm set iDRAC.IPv4.DNSFromDHCP 0
racadm set iDRAC.IPv4.DNSFromDHCP.DNS1 <primary DNS IP address>
racadm set iDRAC.IPv4.DNSFromDHCP.DNS2 <secondary DNS IP address>
```

4. If you want to configure a list of user domains so that you only need to enter the user name during log in to iDRAC web interface, use the following command:

```
racadm set iDRAC.UserDomain.<index>.Name <fully qualified domain name or IP Address
of the domain controller>
```

You can configure up to 40 user domains with index numbers between 1 and 40.

Testing Active Directory settings

You can test the Active Directory settings to verify whether your configuration is correct, or to diagnose the problem with a failed Active Directory log in.

Testing Active Directory settings using iDRAC web interface

To test the Active Directory settings:

- In iDRAC Web Interface, go to iDRAC Settings > Users > Directory Services > Microsoft Active Directory, click Test. The Test Active Directory Settings page is displayed.
- 2. Click Test.
- Enter a test user's name (for example, username@domain.com) and password and click Start Test. A detailed test results and the test log displays.

If there is a failure in any step, examine the details in the test log to identify the problem and a possible solution.

NOTE: When testing Active Directory settings with Enable Certificate Validation checked, iDRAC requires that the Active Directory server be identified by the FQDN and not an IP address. If the Active Directory server is identified by an IP address, certificate validation fails because iDRAC is not able to communicate with the Active Directory server.

Testing Active Directory settings using RACADM

To test the Active Directory settings, use the testfeature command.

For more information, see the iDRAC RACADM CLI Guide available at https://www.dell.com/idracmanuals.

Configuring generic LDAP users

iDRAC provides a generic solution to support Lightweight Directory Access Protocol (LDAP)-based authentication. This feature does not require any schema extension on your directory services.

To make iDRAC LDAP implementation generic, the commonality between different directory services is utilized to group users and then map the user-group relationship. The directory service specific action is the schema. For example, they may have different attribute names for the group, user, and the link between the user and the group. These actions can be configured in iDRAC.

NOTE: The Smart Card based Two Factor Authentication (TFA) and the Single Sign-On (SSO) logins are not supported for generic LDAP Directory Service.



Configuring generic LDAP directory service using iDRAC webbased interface

To configure the generic LDAP directory service using Web interface:

(i) NOTE: For information about the various fields, see the *iDRAC Online Help*.

- In the iDRAC Web interface, go to iDRAC Settings > Users > Directory Services > Generic LDAP Directory Service, click Edit. The Generic LDAP Configuration and Management Step 1 of 3 page displays the current generic LDAP settings.
- Optionally, enable certificate validation and upload the digital certificate used during initiation of SSL connections when communicating with a generic LDAP server.

(i) NOTE: In this release, non-SSL port based LDAP bind is not supported. Only LDAP over SSL is supported.

3. Click Next.

- The Generic LDAP Configuration and Management Step 2 of 3 page is displayed.
- 4. Enable generic LDAP authentication and specify the location information about generic LDAP servers and user accounts.
 - (i) **NOTE:** If certificate validation is enabled, specify the LDAP Server's FQDN and make sure that DNS is configured correctly under **iDRAC Settings** > **Network**.
 - (i) **NOTE:** In this release, nested group is not supported. The firmware searches for the direct member of the group to match the user DN. Also, only single domain is supported. Cross domain is not supported.
- 5. Click Next.
 - The Generic LDAP Configuration and Management Step 3a of 3 page is displayed.
- 6. Click Role Group.
- The Generic LDAP Configuration and Management Step 3b of 3 page is displayed.
- 7. Specify the group distinguished name, the privileges associated with the group, and click **Apply**.
 - (i) **NOTE:** If you are using Novell eDirectory and if you have used these characters—#(hash), "(double quotes), ;(semi colon), > (greater than), , (comma), or <(lesser than)—for the Group DN name, they must be escaped.

The role group settings are saved. The **Generic LDAP Configuration and Management Step 3a of 3** page displays the role group settings.

- 8. If you want to configure additional role groups, repeat steps 7 and 8.
- 9. Click Finish. The generic LDAP directory service is configured.

Configuring generic LDAP directory service using RACADM

To configure the LDAP directory service, use the objects in the iDRAC.LDAP and iDRAC.LDAPRole groups.

For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Testing LDAP directory service settings

You can test the LDAP directory service settings to verify whether your configuration is correct, or to diagnose the problem with a failed LDAP log in.

Testing LDAP directory service settings using iDRAC web interface

To test the LDAP directory service settings:

- 1. In iDRAC Web Interface, go to **iDRAC Settings** > **Users** > **Directory Services** > **Generic LDAP Directory Service**. The **Generic LDAP Configuration and Management** page displays the current generic LDAP settings.
- 2. Click Test.
- **3.** Enter the user name and password of a directory user that is chosen to test the LDAP settings. The format depends on the *Attribute of User Login* is used and the user name entered must match the value of the chosen attribute.



- **NOTE:** When testing LDAP settings with **Enable Certificate Validation** checked, iDRAC requires that the LDAP server be identified by the FQDN and not an IP address. If the LDAP server is identified by an IP address, certificate validation fails because iDRAC is not able to communicate with the LDAP server.
- **NOTE:** When generic LDAP is enabled, iDRAC first tries to login the user as a directory user. If it fails, local user lookup is enabled.

The test results and the test log are displayed.

Testing LDAP directory service settings using RACADM

To test the LDAP directory service settings, use the testfeature command. For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.



System Configuration Lockdown mode

System Configuration Lockdown mode helps in preventing unintended changes after a system is provisioned. Lockdown mode is applicable to both configuration and firmware updates. When the system is locked down, any attempt to change the system configuration is blocked. If any attempts are made to change the critical system settings, an error message is displayed. Enabling System lockdown mode blocks the firmware update of third party I/O cards using the vendor tools.

System Lockdown mode is only available for Enterprise licensed customers.

In 4.40.00.00 release, System lockdown functionality is extended to NIC's also.

- (i) NOTE: Enhanced Lockdown for NIC's only includes firmware lockdown to prevent firmware updates. Configuration (x-UEFI) lockdown is not supported.
- () NOTE: After the System Lockdown mode is enabled, you cannot change any configuration settings. System settings fields are disabled.

Lockdown mode can be enabled or disabled using the following interfaces:

- iDRAC web interface
- RACADM
- WSMan
- SCP (System Configuration Profile)
- Redfish
- Using F2 during POST and selecting iDRAC Settings
- Factory System Erase

NOTE: To enable Lockdown mode, you must have iDRAC Enterprise or Datacenter license and Control and Configure system privileges.

NOTE: You may be able to access vMedia while system is in Lockdown mode but configuring remote file share is not enabled.

(i) NOTE: The interfaces like OMSA, SysCfg, and USC can only check the settings but cannot modify the configurations.

The following table lists the functional and nonfunctional features, interfaces, and utilities that are affected by Lockdown mode: **NOTE:** Changing the boot order using iDRAC is not supported when Lockdown mode is enabled. However, boot-control option is available in vConsole menu, which has no effect when iDRAC is in Lockdown mode.

Table 32. Items affected by Lockdown mode

Disabled	Remains functional
Deleting Licenses	Power Operations - Power ON/OFF, Reset
DUP updates	Power cap setting
SCP import	Power priority
Reset to defaults	Identify devices (Chassis or PERC)
OMSA/OMSS	• Part replacement, Easy Restore, and system board replacement
• IPMI	Running diagnostics
• DRAC/LC	Modular operations (FlexAddress or Remote-Assigned Address)
• DTK-Syscfg	Group Manager passcode
Redfish	• All vendor tools that have direct access to the device (excludes
OpenManage Essentials	selected NIC's)
 BIOS (F2 settings become read-only) 	License export
Group manager	• PERC
Select network cards	PERC CLI
	 DTK-RAIDCFG
	• F2/Ctrl+R



Table 32. Items affected by Lockdown mode

Disabled	Remains functional
	 All Vendor tools that have direct access to the device NVMe DTK-RAIDCFG F2/Ctrl+R BOSS-S1 Marvell CLI F2/Ctrl+R ISM/OMSA settings (OS BMC enable, watchdog ping, OS name, OS version)

(i) NOTE: When lockdown mode is enabled, OpenID Connect login option is not displayed in iDRAC login page.

System Configuration Lockdown mode 165



Configuring iDRAC for Single Sign-On or smart card login

This section provides information to configure iDRAC for Smart Card login (for local users and Active Directory users), and Single Sign-On (SSO) login (for Active Directory users.) SSO and smart card login are licensed features.

iDRAC supports Kerberos based Active Directory authentication to support Smart Card and SSO logins. For information on Kerberos, see the Microsoft website.

Topics:

- Prerequisites for Active Directory Single Sign-On or smart card login
- Configuring iDRAC SSO login for Active Directory users
- Enabling or disabling smart card login
- Configuring Smart Card Login
- Using Smart Card to Login

Prerequisites for Active Directory Single Sign-On or smart card login

The prerequisites to Active Directory based SSO or Smart Card logins are:

- Synchronize iDRAC time with the Active Directory domain controller time. If not, kerberos authentication on iDRAC fails. You can use the Time zone and NTP feature to synchronize the time. To do this, see Configuring time zone and NTP.
- Register iDRAC as a computer in the Active Directory root domain.
- Generate a keytab file using the ktpass tool.
- To enable Single Sign-On for Extended schema, make sure that the Trust this user for delegation to any service (Kerberos only) option is selected on the Delegation tab for the keytab user. This tab is available only after creating the keytab file using ktpass utility.
- Configure the browser to enable SSO login.
- Create the Active Directory objects and provide the required privileges.
- For SSO, configure the reverse lookup zone on the DNS servers for the subnet where iDRAC resides.
 - (i) NOTE: If the host name does not match the reverse DNS lookup, Kerberos authentication fails.
- Configure the browser to support SSO login. For more information, see Single Sign-On.

(i) NOTE: Google Chrome and Safari do not support Active Directory for SSO login.

Registering iDRAC on Domain name System

To register iDRAC in Active Directory root domain:

- 1. Click **iDRAC Settings** > **Connectivity** > **Network**. The **Network** page is displayed.
- 2. You can select IPv4 Settings or IPv6 Settings based on the IP settings.
- **3.** Provide a valid **Preferred/Alternate DNS Server** IP address. This value is a valid DNS server IP address that is part of the root domain.
- 4. Select Register iDRAC on DNS.
- 5. Provide a valid DNS Domain Name.
- **6.** Verify that network DNS configuration matches with the Active Directory DNS information. For more information about the options, see the *iDRAC Online Help*.

166 Configuring iDRAC for Single Sign-On or smart card login



Creating Active Directory objects and providing privileges

Logging in to Active Directory Standard schema based SSO

Perform the following steps for Active Directory Standard schema based SSO login:

- **1.** Create a User Group.
- **2.** Create a User for Standard schema.
- **NOTE:** Use the existing AD User Group & AD User.

Logging in to Active Directory Extended schema based SSO

Perform the following steps for Active Directory Extended schema based SSO login:

- 1. Create the device object, privilege object, and association object in the Active Directory server.
- 2. Set access privileges to the created privilege object.
 - (i) NOTE: It is recommended not to provide administrator privileges as this could bypass some security checks.
- **3.** Associate the device object and privilege object using the association object.
- 4. Add the preceding SSO user (login user) to the device object.
- 5. Provide access privilege to Authenticated Users for accessing the created association object.

Logging in to Active Directory SSO

Perform the following steps for Active Directory SSO login:

1. Create a Kerberos key-tab user which is used for the creation of the key-tab file.

(i) NOTE: Create new KERBROS key for every iDRAC IP.

Configuring iDRAC SSO login for Active Directory users

Before configuring iDRAC for Active Directory SSO login, make sure that you have completed all the prerequisites. You can configure iDRAC for Active Directory SSO when you setup an user account based on Active Directory.

Creating a User in Active Directory for SSO

To create a user in Active Directory for SSO:

- 1. Create a new user in the organization unit.
- 2. Go to Kerberos User>Properties>Account>Use Kerberos AES Encryption types for this account
- **3.** Use the following command to generate a Kerberos keytab in the Active Directory server:

```
C:\> ktpass.exe -princ HTTP/idrac7name.domainname.com@DOMAINNAME.COM -mapuser
DOMAINNAME\username -mapop set -crypto AES256-SHA1 -ptype KRB5_NT_PRINCIPAL -pass
[password] -out c:\krbkeytab
```

Note for Extended Schema

- Change the Delegation setting of the Kerberos user.
- Go to Kerberos User>Properties>Delegation>Trust this user for delegation to any service (Kerberos only)

(i) NOTE: Log-off and Log-in from the Management Station Active Directory user after changing the above setting.

Configuring iDRAC for Single Sign-On or smart card login 167

Inserido ao protocolo **22.951.206-4** por: **Pedro Henrique Golin Linhares** em: 31/03/2025 12:49. A autenticidade deste documento pode ser validada no endereço: https://www.eprotocolo.pr.gov.br/spiweb/validarDocumento com o código: 692a5ef82beb173ed33dad4ee6505546.



Generating Kerberos keytab file

To support the SSO and smart card login authentication, iDRAC supports the configuration to enable itself as a kerberized service on a Windows Kerberos network. The Kerberos configuration on iDRAC involves the same steps as configuring a non–Windows Server Kerberos service as a security principal in Windows Server Active Directory.

The *ktpass* tool (available from Microsoft as part of the server installation CD/DVD) is used to create the Service Principal Name (SPN) bindings to a user account and export the trust information into a MIT-style Kerberos *keytab* file, which enables a trust relation between an external user or system and the Key Distribution Centre (KDC). The keytab file contains a cryptographic key, which is used to encrypt the information between the server and the KDC. The ktpass tool allows UNIX-based services that support Kerberos authentication to use the interoperability features provided by a Windows Server Kerberos KDC service. For more information on the **ktpass** utility, see the Microsoft website at: **technet.microsoft.com/en-us/library/cc779157(WS.10).aspx**

Before generating a keytab file, you must create an Active Directory user account for use with the **-mapuser** option of the ktpass command. Also, you must have the same name as iDRAC DNS name to which you upload the generated keytab file.

To generate a keytab file using the ktpass tool:

- 1. Run the *ktpass* utility on the domain controller (Active Directory server) where you want to map iDRAC to a user account in Active Directory.
- 2. Use the following ktpass command to create the Kerberos keytab file:

```
C:\> ktpass.exe -princ HTTP/idrac7name.domainname.com@DOMAINNAME.COM -mapuser
DOMAINNAME\username -mapop set -crypto AES256-SHA1 -ptype KRB5_NT_PRINCIPAL -pass
[password] -out c:\krbkeytab
```

The encryption type is AES256-SHA1. The principal type is KRB5_NT_PRINCIPAL. The properties of the user account to which the Service Principal Name is mapped to must have **Use AES 256 encryption types for this account** property enabled.

NOTE: Use lowercase letters for the **iDRACname** and **Service Principal Name**. Use uppercase letters for the domain name as shown in the example.

A keytab file is generated.

NOTE: If you find any issues with iDRAC user for which the keytab file is created, create a new user and a new keytab file. If the same keytab file which was initially created is again executed, it does not configure correctly.

Configuring iDRAC SSO login for Active Directory users using web interface

To configure iDRAC for Active Directory SSO login:

(i) NOTE: For information about the options, see the *iDRAC Online Help*.

- 1. Verify whether the iDRAC DNS name matches the iDRAC Fully Qualified Domain Name. To do this, in iDRAC Web interface, go to **iDRAC Settings** > **Network** > **Common Settings** and refer to **DNS iDRAC Name** property.
- 2. While configuring Active Directory to setup a user account based on standard schema or extended schema, perform the following two additional steps to configure SSO:
 - Upload the keytab file on the Active Directory Configuration and Management Step 1 of 4 page.
 - Select Enable Single Sign-On option on the Active Directory Configuration and Management Step 2 of 4 page.

Configuring iDRAC SSO login for Active Directory users using RACADM

To enable SSO, complete the steps to configure Active Directory, and run the following command:

racadm set iDRAC.ActiveDirectory.SSOEnable 1

168 Configuring iDRAC for Single Sign-On or smart card login



Management Station Settings

Perform the following steps after configuring SSO login for Active Directory users:

- 1. Set the DNS Server IP in Network properties and mention the preferred DNS Server IP.
- 2. Go to My Computer and add the *domain.tld domain.
- 3. Add the Active Directory User to Administrator by navigating to: My Computer > Manage > Local User and Groups >
- **Groups** > **Administrator** and add the Active Directory User.
- Logoff the system and login using the Active Directory User credential.
 In Internet Explorer Setting, add *domain.tld domain as below:
 - a. Go to Tools > Internet Options > Security > Local Internet > Sites and clear the Automatically detect intranet network setting selection. Select the remaining three options and click Advanced to add *domain.tld.
 - b. Open a new window in IE and use the iDRAC hostname to launch the iDRAC GUI.
- 6. In Mozilla Firefox Setting, add *domain.tld domain:
 - Launch Firefox browser and type about:config in the URL.
 - Use negotiate in the filter textbox. Double click the result consisting of *auth.trusted.uris*. Type the domain, save the settings and close the browser.
 - Open a new window in Firefox and use the iDRAC hostname to launch the iDRAC GUI.

Enabling or disabling smart card login

Before enabling or disabling smart card login for iDRAC, make sure that:

- You have configure iDRAC permissions.
- iDRAC local user configuration or Active Directory user configuration with the appropriate certificates is complete.

NOTE: If smart card login is enabled, then SSH, IPMI Over LAN, Serial Over LAN, and remote RACADM are disabled. Again, if you disable smart card login, the interfaces are not enabled automatically.

Enabling or disabling smart card login using web interface

To enable or disable the Smart Card logon feature:

- In the iDRAC web interface, go to iDRAC Settings > Users > Smart Card. The Smart Card page is displayed.
- 2. From the **Configure Smart Card Logon** drop-down menu, select **Enabled** to enable smart card logon or select **Enabled With Remote RACADM**. Else, select **Disabled**.

For more information about the options, see the *iDRAC Online Help*.

 Click Apply to apply the settings. You are prompted for a Smart Card login during any subsequent logon attempts using the iDRAC web interface.

Enabling or disabling smart card login using RACADM

To enable smart card login, use the set command with objects in the iDRAC.SmartCard group.

For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Enabling or disabling smart card login using iDRAC settings utility

To enable or disable the Smart Card logon feature:

- 1. In the iDRAC Settings utility, go to Smart Card. The **iDRAC Settings Smart Card** page is displayed.
- 2. Select **Enabled** to enable smart card logon. Else, select **Disabled**. For more information about the options, see *iDRAC Settings Utility Online Help*.
- **3.** Click **Back**, click **Finish**, and then click **Yes**. The smart card logon feature is enabled or disabled based on the selection.



Configuring Smart Card Login

(i) **NOTE:** For Active Directory Smart Card Configuration, iDRAC must be configured either with Standard or Extended Schema SSO Login.

Configuring iDRAC smart card login for Active Directory users

Before configuring iDRAC Smart Card login for Active Directory users, make sure that you have completed the required prerequisites.

To configure iDRAC for smart card login:

- 1. In iDRAC Web interface, while configuring Active Directory to set up an user account based on standard schema or extended schema, on the Active Directory Configuration and Management Step 1 of 4 page:
 - Enable certificate validation.
 - Upload a trusted CA-signed certificate.
 - Upload the keytab file.
- 2. Enable smart card login. For information about the options, see the *iDRAC Online Help*.

Configuring iDRAC smart card login for local users

To configure iDRAC local user for smart card login:

- 1. Upload the smart card user certificate and trusted CA certificate to iDRAC.
- 2. Enable smart card login.

Uploading smart card user certificate

Before you upload the user certificate, make sure that the user certificate from the smart card vendor is exported in Base64 format. SHA-2 certificates are also supported.

Uploading smart card user certificate using web interface

To upload smart card user certificate:

1. In iDRAC web interface, go to iDRAC Settings > Users > Smart Card.

(i) NOTE: The Smart Card login feature requires the configuration of the local and/or Active Directory user certificate.

- 2. Under Configure Smart Card Logon, select Enabled With Remote RACADM to enable the configuration..
- 3. Set the option to Enable CRL Check for Smart Card Logon.
- 4. Click Apply.

Uploading smart card user certificate using RACADM

To upload smart card user certificate, use the **usercertupload** object. For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Requesting Certificate for smart card enrollment

Follow these steps to request certificate for smart card enrollment:

- 1. Connect the smart card in the client system and install the required drivers & software.
- 2. Verify the driver status in the Device Manager.
- 3. Launch the smart card enrollment agent in the browser.
- 4. Enter the Username & Password and click OK.
- 5. Click Request Certificate.

170 Configuring iDRAC for Single Sign-On or smart card login



6. Click Advanced Certificate Request.

- 7. Click **Request a certificate** for a smart card on behalf of another user by using the smart card certificate enrollment station.
- 8. Select user to enroll by clicking **Select User** button.
- 9. Click **Enroll** and enter the smart card credential.
- 10. Enter the smart card PIN and click on **Submit**.

Uploading trusted CA certificate for smart card

Before you upload the CA certificate, make sure that you have a CA-signed certificate.

Uploading trusted CA certificate for smart card using web interface

To upload trusted CA certificate for smart card login:

- In iDRAC Web interface, go to iDRAC Settings > Network > User Authentication > Local Users. The Users page is displayed.
- 2. In the User ID column, click a user ID number. The Users Main Menu page is displayed.
- **3.** Under Smart Card Configurations, select Upload Trusted CA Certificate and click Next. The Trusted CA Certificate Upload page is displayed.
- 4. Browse and select the trusted CA certificate, and click **Apply**.

Uploading trusted CA certificate for smart card using RACADM

To upload trusted CA certificate for smart card login, use the **usercertupload** object. For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Using Smart Card to Login

(i) NOTE: Smart card login is supported only on Internet Explorer.

To login using smart card:

- 1. Logout from iDRAC GUI after enabling smart card.
- 2. Launch iDRAC by using http://IP/ or launch using FQDN http://FQDN/
- 3. Click Install after smart card plug-in download.
- 4. Enter smart card PIN and click Submit.
- 5. iDRAC will login successfully using smart card.



Configuring iDRAC to send alerts

You can set alerts and actions for certain events that occur on the managed system. An event occurs when the status of a system component is greater than the pre-defined condition. If an event matches an event filter and you have configured this filter to generate an alert (e-mail, SNMP trap, IPMI alert, remote system logs, Redfish event, or WS events), then an alert is sent to one or more configured destinations. If the same event filter is also configured to perform an action (such as reboot, power cycle, or power off the system), the action is performed. You can set only one action for each event.

To configure iDRAC to send alerts:

- 1. Enable alerts.
- 2. Optionally, you can filter the alerts based on category or severity.
- **3.** Configure the e-mail alert, IPMI alert, SNMP trap, remote system log, Redfish event, operating system log, and/or WS-event settings.
- 4. Enable event alerts and actions such as:
 - Send an email alert, IPMI alert, SNMP traps, remote system logs, Redfish event, operating system log, or WS events to configured destinations.
 - Perform a reboot, power off, or power cycle the managed system.

Topics:

- Enabling or disabling alerts
- Filtering alerts
- Setting event alerts
- Setting alert recurrence event
- Setting event actions
- Configuring email alert, SNMP trap, or IPMI trap settings
- Configuring WS Eventing
- Configuring Redfish Eventing
- Monitoring chassis events
- Alerts message IDs

Enabling or disabling alerts

For sending an alert to configured destinations or to perform an event action, you must enable the global alerting option. This property overrides individual alerting or event actions that is set.

Enabling or disabling alerts using web interface

To enable or disable generating alerts:

- In iDRAC web interface, go to Configuration > System Settings > Alert Configuration. The Alerts page is displayed.
- 2. Under Alerts section:
 - Select **Enable** to enable alert generation or perform an event action.
 - Select **Disable** to disable alert generation or disable an event action.
- 3. Click Apply to save the setting.

Quick Alert Configuration

To configure alerts in bulk:

1. Go to Quick Alert Configuration under Alert Configuration page.

172 Configuring iDRAC to send alerts


2. Under Quick Alert Configuration section:

- Select the alert category.
- Select the issue severity notification.
- Select the location where you would like to receive these notifications.
- 3. Click Apply to save the setting.

(i) NOTE: You must select at least one category, one severity, and one destination type to apply the configuration.

All the alerts that are configured are displayed in total under **Alerts Configuration Summary**.

Enabling or disabling alerts using RACADM

Use the following command:

racadm set iDRAC.IPMILan.AlertEnable <n>

```
n=0 — Disabled
```

n=1 — Enabled

Enabling or disabling alerts using iDRAC settings utility

To enable or disable generating alerts or event actions:

- In the iDRAC Settings utility, go to Alerts. The iDRAC Settings Alerts page is displayed.
- 2. Under **Platform Events**, select **Enabled** to enable alert generation or event action. Else, select **Disabled**. For more information about the options, see *iDRAC Settings Utility Online Help*.
- **3.** Click **Back**, click **Finish**, and then click **Yes**. The alert settings are configured.

Filtering alerts

You can filter alerts based on category and severity.

Filtering alerts using iDRAC web interface

To filter the alerts based on category and severity:

(i) NOTE: Even if you are a user with read-only privileges, you can filter the alerts.

- 1. In iDRAC Web interface, go to Configuration > System Settings > Alerts and Remote System Log Configuration.
- 2. Under Alerts and Remote System Log Configuration section, select Filter:
 - System Health System Health category represents all the alerts that are related to hardware within the system chassis. Examples include temperature errors, voltage errors, device errors.
 - Storage Health Storage Health category represents alerts that are related to the storage subsystem. Examples include, controller errors, physical disk errors, virtual disk errors.
 - Configuration Configuration category represents alerts that are related to hardware, firmware and software configuration changes. Examples include, PCI-e card added/removed, RAID configuration changed, iDRAC license changed.
 - Audit Audit category represents the audit log. Examples include, user login/logout information, Password authentication failures, session info, power states.
 - Updates Update category represents alerts that are generated due to firmware/Driver upgrades/downgrades.
 NOTE: This doesn't represent firmware inventory.
 - Work Notes
- 3. Select one or more of the following severity levels:



- Informational
- Warning
- Critical
- 4. Click Apply.
 - The Alert Results section displays the results based on the selected category and severity.

Filtering alerts using RACADM

To filter the alerts, use the **eventfilters** command. For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Setting event alerts

You can set event alerts such as e-mail alerts, IPMI alerts, SNMP traps, remote system logs, operating system logs, and WS events to be sent to configured destinations.

Setting event alerts using web interface

To set an event alert using the web interface:

- 1. Make sure that you have configured the e-mail alert, IPMI alert, SNMP trap settings, and/or remote system log settings.
- 2. In iDRAC Web interface, go to Configuration > System Settings > Alerts and Remote System Log Configuration.
- 3. Under Category, select one or all of the following alerts for the required events:
 - Email
 - SNMP Trap
 - IPMI Alert
 - Remote System Log
 - WS Eventing
 - OS Log
 - Redfish Event
- 4. Select Action.
 - The setting is saved.
- 5. Optionally, you can send a test event. In the **Message ID to Test Event** field, enter the message ID to test if the alert is generated and click **Test**. For more information about the event and error messages generated by the system firmware and agents that monitor system components, see the *Event and Error Message Reference Guide* at iDRACmanuals

Setting event alerts using RACADM

To set an event alert, use the **eventfilters** command. For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Setting alert recurrence event

You can configure iDRAC to generate additional events at specific intervals if the system continues to operate at a temperature which is greater than the inlet temperature threshold limit. The default interval is 30 days. The valid range is 0 to 366 days. A value of '0' indicates no event recurrence.

(i) NOTE: You must have Configure iDRAC privilege to set the alert recurrence value.

Setting alert recurrence events using RACADM

To set the alert recurrence event using RACADM, use the **eventfilters** command. For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Inserido ao protocolo **22.951.206-4** por: **Pedro Henrique Golin Linhares** em: 31/03/2025 12:49. A autenticidade deste documento pode ser validada no endereço: https://www.eprotocolo.pr.gov.br/spiweb/validarDocumento com o código: 692a5ef82beb173ed33dad4ee6505546.



Setting alert recurrence events using iDRAC web interface

To set the alert recurrence value:

- 1. In iDRAC Web interface, go to Configuration > System Settings > Alert Recurrence.
- 2. In the **Recurrence** column, enter the alert frequency value for the required category, alert, and severity type(s). For more information, see the *iDRAC Online help*.
- 3. Click Apply.

The alert recurrence settings are saved.

Setting event actions

You can set event actions such as perform a reboot, power cycle, power off, or perform no action on the system.

Setting event actions using web interface

To set an event action:

- 1. In iDRAC Web interface, go to Configuration > System Settings > Alert and Remote System Log Configuration.
- 2. From the Actions drop-down menu, for each event select an action:
 - Reboot
 - Power Cycle
 - Power Off
 - No Action
- Click Apply. The setting is saved.

Setting event actions using RACADM

To configure an event action, use the eventfilters command. For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Configuring email alert, SNMP trap, or IPMI trap settings

The management station uses Simple Network Management Protocol (SNMP) and Intelligent Platform Management Interface (IPMI) traps to receive data from iDRAC. For systems with large number of nodes, it may not be efficient for a management station to poll each iDRAC for every condition that may occur. For example, event traps can help a management station with load balancing between nodes or by issuing an alert if an authentication failure occurs. SNMP v1, v2, and v3 formats are supported.

You can configure the IPv4 and IPv6 alert destinations, email settings, and SMTP server settings, and test these settings. You can also specify the SNMP v3 user to whom you want to send the SNMP traps.

Before configuring the email, SNMP, or IPMI trap settings, make sure that:

- You have Configure RAC permission.
- You have configured the event filters.

Configuring IP alert destinations

You can configure the IPv6 or IPv4 addresses to receive the IPMI alerts or SNMP traps.

For information about the iDRAC MIBs required to monitor the servers using SNMP, see the *Dell EMC OpenManage SNMP Reference Guide* available at https://www.dell.com/openmanagemanuals.

Configuring iDRAC to send alerts 175



Configuring IP alert destinations using web interface

To configure alert destination settings using Web interface:

- 1. In iDRAC Web interface, go to Configuration > System Settings > SNMP and E-mail Settings.
- Select the State option to enable an alert destination (IPv4 address, IPv6 address, or Fully Qualified Domain Name (FQDN)) to receive the traps.

You can specify up to eight destination addresses. For more information about the options, see the *iDRAC Online Help*.

- **3.** Select the SNMP v3 user to whom you want to send the SNMP trap.
- **4.** Enter the iDRAC SNMP community string (applicable only for SNMPv1 and v2) and the SNMP alert port number. For more information about the options, see the *iDRAC Online Help*.
 - () NOTE: The Community String value indicates the community string to use in a Simple Network Management Protocol (SNMP) alert trap sent from iDRAC. Make sure that the destination community string is the same as the iDRAC community string. The default value is Public.
- To test whether the IP address is receiving the IPMI or SNMP traps, click Send under Test IPMI Trap and Test SNMP Trap respectively.
- 6. Click Apply.
 - The alert destinations are configured.
- In the SNMP Trap Format section, select the protocol version to be used to send the traps on the trap destination(s) SNMP v1, SNMP v2, or SNMP v3 and click Apply.

NOTE: The **SNMP Trap Format** option applies only for SNMP Traps and not for IPMI Traps. IPMI Traps are always sent in SNMP v1 format and is not based on the configured **SNMP Trap Format** option.

The SNMP trap format is configured.

Configuring IP alert destinations using RACADM

To configure the trap alert settings:

1. To enable traps:

```
racadm set idrac.SNMP.Alert.<index>.Enable <n>
```

Parameter	Description	
<index></index>	Destination index. Allowed values are 1 through 8.	
<n>=0</n>	Disable the trap	
<n>=1</n>	Enable the trap	

2. To configure the trap destination address:

```
racadm set idrac.SNMP.Alert.<index>.DestAddr <Address>
```

Parameter	Description	
<index> Destination index. Allowed values are 1 through 8.</index>		
<address> A valid IPv4, IPv6, or FQDN address</address>		

3. Configure the SNMP community name string:

racadm set idrac.ipmilan.communityname <community_name>

Parameter	Description	
<community_name></community_name>	The SNMP Community Name.	

4. To configure SNMP destination:

• Set the SNMP trap destination for SNMPv3:



racadm set idrac.SNMP.Alert.<index>.DestAddr <IP address>

• Set SNMPv3 users for trap destinations:

racadm set idrac.SNMP.Alert.<index>.SNMPv3Username <user_name>

• Enable SNMPv3 for a user:

racadm set idrac.users.<index>.SNMPv3Enable Enabled

5. To test the trap, if required:

racadm testtrap -i <index>

For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Configuring IP alert destinations using iDRAC settings utility

You can configure alert destinations (IPv4, IPv6, or FQDN) using the iDRAC Settings utility. To do this:

- In the iDRAC Settings utility, go to Alerts. The iDRAC Settings Alerts page is displayed.
- 2. Under **Trap Settings**, enable the IP address(es) to receive the traps and enter the IPv4, IPv6, or FQDN destination address(es). You can specify up to eight addresses.
- **3.** Enter the community string name. For information about the options, see the *iDRAC Settings Utility Online Help*.
- **4.** Click **Back**, click **Finish**, and then click **Yes**. The alert destinations are configured.

Configuring email alert settings

You can configure the sender email address and receiver (destination) email address to receive the email alerts. Also, configure the SMTP server address settings.

- **NOTE:** Email alerts support both IPv4 and IPv6 addresses. The iDRAC DNS Domain Name must be specified when using IPv6.
- **NOTE:** If you are using an external SMTP server, ensure that iDRAC can communicate with that server. If the server is unreachable, the error RAC0225 is displayed while trying to send a test mail.

Configuring email alert settings using web interface

To configure the email alert settings using Web interface:

- 1. In iDRAC Web interface, go to Configuration > System Settings > SMTP (E-mail) Configuration.
- **2.** Type a valid email address.
- 3. Click Send under Test Email to test the configured email alert settings.
- 4. Click Apply.
- 5. For SMTP (E-mail) Server Settings provide the following details:
 - SMTP (E-mail) Server IP Address or FQDN/DNS Name
 - Custom Sender Address This field has the following options:
 - **Default** Address field is not editable
 - \circ **Custom** You can enter the email ID from which you can receive the email alerts
 - Custom Message Subject Prefix This field has the following options:
 - Default Default message is not editable
 - Custom You can choose the message to appear in the Subject line of the email
 - SMTP Port Number The connection can be encrypted and emails can be sent over secure ports:

Configuring iDRAC to send alerts 177



- No Encryption port 25 (default)
- SSL --- Port 465
- Connection Encryption When you do not have an email server in your premises, you can use cloud based email servers or SMTP Relays. To configure cloud email server, you can set this feature to any of the following values from the drop down:
 - None No encryption on the connection to the SMTP server. It is the default value.
 - **SSL** Runs SMTP protocol over SSL

() NOTE:

- This feature is not configurable via Group Manager.
- This is a licensed feature and is not available in iDRAC Basic License.
- You must have Configure iDARC privilege to use this feature.
- Authentication
- Username

For Server settings, the port usage depends on connectionencryptiontype and this can be configured only using RACADM.

6. Click Apply. For more information about the options, see the *iDRAC Online Help*.

Configuring email alert settings using RACADM

1. To enable email alert:

```
racadm set iDRAC.EmailAlert.Enable.[index] [n]
```

Parameter	Description
index	Email destination index. Allowed values are 1 through 4.
n=0	Disables email alerts.
n=1	Enables email alerts.

2. To configure email settings:

```
racadm set iDRAC.EmailAlert.Address.[index] [email-address]
```

Parameter	Description
index	Email destination index. Allowed values are 1 through 4.
email-address	Destination email address that receives the platform event alerts.

3. To configure sender email settings:

```
racadm set iDRAC.RemoteHosts.[index] [email-address]
```

Parameter	Description	
index	Sender Email index.	
email-address Sender email address that sends the platform event alerts.		

4. To configure a custom message:

```
racadm set iDRAC.EmailAlert.CustomMsg.[index] [custom-message]
```

Parameter	Description
index	Email destination index. Allowed values are 1 through 4.

/	220TOCOL	\geq
NIN 1	FIs. <u>565</u> Mov. <u>32</u>	BB
	HADO DO E	/

Parameter	Description
custom-message (Custom message

5. To test the configured email alert, if required:

racadm testemail -i [index]

Parameter	Description
index	Email destination index to be tested. Allowed values are 1 through 4.

For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Configuring SMTP email server address settings

You must configure the SMTP server address for email alerts to be sent to specified destinations.

Configuring SMTP email server address settings using iDRAC web interface

To configure the SMTP server address:

- 1. In iDRAC Web interface, go to Configuration > System Settings > Alert Configuration > SNMP (E-mail Configuration).
- 2. Enter the valid IP address or fully qualified domain name (FQDN) of the SMTP server to be used in the configuration.
- **3.** Select the **Enable Authentication** option and then provide the user name and password (of a user who has access to SMTP server).
- Enter the SMTP port number.
 For more information about the fields, see the *iDRAC Online Help*.
- 5. Click **Apply**. The SMTP settings are configured.

Configuring SMTP email server address settings using RACADM

To configure the SMTP email server:

racadm set iDRAC.RemoteHosts.SMTPServerIPAddress <SMTP E-mail Server IP Address>

Configuring WS Eventing

The WS Eventing protocol is used for a client service (subscriber) to register interest (subscription) with a server (event source) for receiving messages containing the server events (notifications or event messages). Clients interested in receiving the WS Eventing messages can subscribe with iDRAC and receive Lifecycle Controller job related events.

The steps required to configure WS Eventing feature to receive WS Eventing messages for changes related to Lifecycle Controller jobs are described in the Web service Eventing Support for iDRAC 1.30.30 specification document. In addition to this specification, see the DSP0226 (DMTF WS Management Specification), Section 10 Notifications (Eventing) document for the complete information on the WS Eventing protocol. The Lifecycle Controller related jobs are described in the DCIM Job Control Profile document.

Configuring Redfish Eventing

The Redfish eventing protocol is used for a client service (subscriber) to register interest (subscription) with a server (event source) for receiving messages containing the Redfish events (notifications or event messages). Clients interested in receiving the Redfish eventing messages can subscribe with iDRAC and receive Lifecycle Controller job related events.



Monitoring chassis events

On the PowerEdge FX2/FX2s chassis, you can enable the **Chassis Management and Monitoring** setting in iDRAC to perform chassis management and monitoring tasks such as monitoring chassis components, configuring alerts, using iDRAC RACADM to pass CMC RACADM commands, and updating the chassis management firmware. This setting allows you to manage the servers in the chassis even if the CMC is not on the network. You can set the value to **Disabled** to forward the chassis events. By default, this setting is set as **Enabled**.

(i) NOTE: For this setting to take effect, you must ensure that in CMC, the Chassis Management at Server setting must be set to Monitor or Manage and Monitor.

When the **Chassis Management and Monitoring** option is set to **Enabled**, iDRAC generates and logs chassis events. The events generated are integrated into the iDRAC event subsystem and alerts are generated similar to the rest of the events.

CMC also forwards the events generated to iDRAC. In case the iDRAC on the server is not functional, CMC queues the first 16 events and logs the rest in the CMC log. These 16 events are sent to iDRAC as soon as **Chassis monitoring** is set to enabled.

In instances where iDRAC detects that a required CMC functionality is absent, a warning message is displayed informing you that certain features may not be functional without a CMC firmware upgrade.

() NOTE: iDRAC doesnot support the following Chassis attributes:

- ChassisBoardPartNumber
- ChassisBoardSerialNumber

Monitoring chassis events using the iDRAC web interface

To monitor chassis events using the iDRAC web interface, perform the following steps:

NOTE: This section appears only for PowerEdge FX2/FX2s chassis and if **Chassis Management at Server** mode is set to **Monitor** or **Manage and Monitor** in CMC.

- 1. On the CMC interface, click Chassis Overview > Setup > General.
- 2. From the Chassis Management at Server Mode drop-down menu, select Manage and Monitor, and click Apply.
- 3. Launch the iDRAC web interface, click **Overview** > iDRAC Settings > CMC.
- 4. Under the Chassis Management at Server section, ensure that Capability from iDRAC drop-down box is set to Enabled.

Monitoring chassis events using RACADM

This setting is applicable only for PowerEdge FX2/FX2s servers and if **Chassis Management at Server** mode is set to **Monitor** or **Manage and Monitor** in CMC.

To monitor chassis events using iDRAC RACADM:

racadm get system.chassiscontrol.chassismanagementmonitoring

For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Alerts message IDs

The following table provides the list of message IDs that are displayed for the alerts.

Table 33. Alert message IDs

Message ID	Description	Description (For MX platforms)
АМР	Amperage	Amperage
ASR	Auto Sys Reset	Auto Sys Reset



Table 33. Alert message IDs (continued)

Message ID	Description	Description (For MX platforms)
ВАТ	Battery Event	Battery Event
BIOS	BIOS Management	BIOS Management
воот	BOOT Control	BOOT Control
CBL	Cable	Cable
CPU	Processor	Processor
CPUA	Proc Absent	Proc Absent
CTL	Storage Contr	Storage Contr
DH	Cert Mgmt	Cert Mgmt
DIS	Auto-Discovery	Auto-Discovery
ENC	Storage Enclosr	Storage Enclosr
FAN	Fan Event	Fan Event
FSD	Debug	Debug
HWC	Hardware Config	Hardware Config
IPA	DRAC IP Change	DRAC IP Change
ITR	Intrusion	Intrusion
JCP	Job Control	Job Control
LC	Lifecycle Controller	Lifecycle Controller
LIC	Licensing	Licensing
LNK	Link Status	Link Status
LOG	Log event	Log event
MEM	Memory	Memory
NDR	NIC OS Driver	NIC OS Driver
NIC	NIC Config	NIC Config
OSD	OS Deployment	OS Deployment
OSE	OS Event	OS Event
PCI	PCI Device	PCI Device
PDR	Physical Disk	Physical Disk
PR	Part Exchange	Part Exchange
PST	BIOS POST	BIOS POST



Table 33. Alert message IDs (continued)

Message ID	Description	Description (For MX platforms)
PSU	Power Supply	Power Supply
PSUA	PSU Absent	PSU Absent
PWR	Power Usage	Power Usage
RAC	RAC Event	RAC Event
RDU	Redundancy	Redundancy
RED	FW Download	FW Download
RFL	IDSDM Media	IDSDM Media
RFLA	IDSDM Absent	IDSDM Absent
RFM	FlexAddress SD	Not Applicable
RRDU	IDSDM Redundancy	IDSDM Redundancy
RSI	Remote Service	Remote Service
SEC	Security Event	Security Event
SEL	Sys Event Log	Sys Event Log
SRD	Software RAID	Software RAID
SSD	PCIe SSD	PCIe SSD
STOR	Storage	Storage
SUP	FW Update Job	FW Update Job
SWC	Software Config	Software Config
SWU	Software Change	Software Change
SYS	System Info	System Info
ТМР	Temperature	Temperature
TST	Test Alert	Test Alert
UEFI	UEFI Event	UEFI Event
USR	User Tracking	User Tracking
VDR	Virtual Disk	Virtual Disk
VF	vFlash SD card	vFlash SD card
VFL	vFlash Event	vFlash Event
VFLA	vFlash Absent	vFlash Absent
VLT	Voltage	Voltage

182 Configuring iDRAC to send alerts

Inserido ao protocolo **22.951.206-4** por: **Pedro Henrique Golin Linhares** em: 31/03/2025 12:49. A autenticidade deste documento pode ser validada no endereço: https://www.eprotocolo.pr.gov.br/spiweb/validarDocumento com o código: **692a5ef82beb173ed33dad4ee6505546**.



Table 33. Alert message IDs (continued)

Message ID	Description	Description (For MX platforms)
VME	Virtual Media	Virtual Media
VRM	Virtual Console	Virtual Console
WRK	Work Note	Work Note



iDRAC 9 Group Manager

Group Manager enables user to have multiple console experience and offers simplified basic iDRAC management.

iDRAC Group Manager feature is available for Dell's 14th generation servers to offer simplified basic management of iDRACs and associated servers on the local network using the iDRAC GUI. Group Manager allows 1XMany console experience without involving a separate application. It allows the users to view the details of a set of servers by permitting more powerful management than by inspecting servers visually for faults and other manual methods.

Group manager is a licensed feature and part of the Enterprise license. Only iDRAC admin users can access the Group Manager functionality.

(i) NOTE: For better user experience Group Manager supports up to 250 server nodes.

Topics:

- Group Manager
- Summary View
- Network Configuration requirements
- Manage Logins
- Configure Alerts
- Export
- Discovered Servers View
- Jobs View
- Jobs Export
- Group Information Panel
- Group Settings
- Actions on a selected Server
- iDRAC Group Firmware Update

Group Manager

To use **Group Manager** feature, you need to enable the **Group Manager** from iDRAC index page or on the Group Manager Welcome screen. The group manager welcome screen provides options listed in the below table.

Table 34. Options in Group Manager

Option	Description
Join Existing Group	 Allows you to join an existing group, you need to know the GroupName and Passcode to join a specific group. (i) NOTE: Passwords are associated to iDRAC user credentials. Whereas, a passcode is associated to a group to establish authenticated device communication between different iDRACs in the same group.
Create New Group	Allows you to create a new group. The specific iDRAC which has created the group would be the master (primary controller) of the group.
Disable Group Manager for this System	You can select this option in case you do not want to join any group from a specific system. However, you can access Group Manager at any point of time by selecting Open Group Manager from the iDRAC index page. Once you disable the group manager, user needs to wait for 60 seconds to perform any further group manager operations.

Inserido ao protocolo **22.951.206-4** por: **Pedro Henrique Golin Linhares** em: 31/03/2025 12:49. A autenticidade deste documento pode ser validada no endereço: https://www.eprotocolo.pr.gov.br/spiweb/validarDocumento com o código: 692a5ef82beb173ed33dad4ee6505546.



Once the group manager feature is enabled, that iDRAC allows you the option to create or join an iDRAC local group. More than one iDRAC group can be setup in the local network but an individual iDRAC can only be a member of one group at a time. To change group (join a new group) the iDRAC must first leave its current group and then join the new group. The iDRAC from where the group was created gets chosen as the primary controller of the group by default. The user does not define a dedicated group manager primary controller to control that group. The primary controller hosts the group manager web interface and provides the GUI based work flows. The iDRAC members self-select a new primary controller for the group if the current primary goes offline for a prolonged duration, but that does not have any impact on the end user. You can normally access the group manager from all iDRAC members by clicking group manager from the iDRAC index page.

Summary View

You need to have administrator privileges to access group manager pages. If a non-administrator user logs onto the iDRAC, the group manager section does not appear with their credentials. The group manager home page (summary view) is broadly categorized as three sections. The first section shows rollup summary with aggregated summary details.

- Total number of servers in the local group.
- Chart showing number of servers per server model.
- Doughnut chart showing the servers per their health status (clicking on a chart section filters the server list to show only the servers with selected health).
- Warning box if there is a duplicate group detected in the local network. Duplicate group is typically the group with the same name but with a different passcode. This warning box does not appear if there is no duplicate group.
- Displays the iDRACs, that are controlling the group (Primary and Secondary controller).

The second section provides buttons for actions that are taken on the group as whole and the third section displays the list of all iDRACs in the group.

It shows all the systems in the group and their current health status and allows the user to take corrective action as needed. Server attributes specific to a server is described in the below table.

Table 35. Server Attributes

Server Attribute	Description
Health	Indicates the health status of that specific server.
Host Name	Displays the Server name.
iDRAC IP Address	Displays the exact IPV4 and IPV6 addresses.
Service Tag	Displays the Service Tag information.
Model	Displays the Model number of the Dell Server.
idrac	Displays the iDRAC version.
Last Status Update	Displays the time stamp when the server Status was last updated.

The System Information panel provides further details on the server like iDRAC network connectivity status, server host power state, express service code, operating system, asset tag, node ID, IDRAC DNS name, Server BIOS version, Server CPU information, System memory and location information. You may double click on a row or click on the launch iDRAC button to perform a single sign on redirect to the selected iDRAC index page. On the selected server, virtual console can be accessed or server power actions can be performed from More Actions dropped down list.

Manage iDRAC user logins, Alert configuration and group inventory export are the group actions supported.

Network Configuration requirements

Group Manager uses IPv6 link local networking to communicate between iDRACs (excluding the web browser GUI). Link local communication is defined as non-routed packets which means any iDRAC separated by a router cannot be joined in a local group. If the iDRAC-dedicated port or shared LOM is assigned to a vLAN, the vLAN limits the number of iDRACs that can be joined in a group (iDRACs must be on same vLAN and traffic must not pass through a router).

When Group Manager is enabled, iDRAC enables an IPv6 Link Local address regardless of the iDRAC's current user defined network configuration. Group Manager can be used when iDRAC is configured for IPv4 or IPv6 IP addresses.



Group Manager uses mDNS to discover other iDRACs on the network and sends encrypted packets for normal inventorying, monitoring and management of the group using the link local IP address. Using IPv6 link local networking means that the Group Manager ports and packets will never leave the local network or be accessible to external networks.

Ports (Specific to Group Manager unique functionality does not include all iDRAC ports) are:

- 5353 (mDNS)
- 443 (webserver) configurable
- 5670 (Multicast group communication)
- C000 -> F000 dynamically identifies one free port for each member to communicate in the group

Best networking practices

- Groups are intended to be small and on the same physical link local network.
- It is recommended to use the dedicated iDRAC network port for enhanced security. Shared LOM is also supported.

Additional network considerations

Two iDRACs that are separated by a router in the network topology are considered to be on separate local networks and cannot be joined in the same iDRAC local group. Meaning, if the iDRAC is configured for dedicated NIC settings, the network cable connected to iDRAC dedicated port in the rear of the server must be under a local network for all relevant servers.

If the iDRAC is configured for shared LOM network settings, the shared network connection used by both server host and IDRAC need to be connected under a local network for Group Manager to detect and onboard those servers into a common group. IDRACs configured with a mix of dedicated and shared LOM mode NIC settings could also be on-boarded into a common group, if all the network connections do not pass through a router.

Effect of MLD snooping in VLAN environments on Group Manager Discovery

Since Group Manager uses IPv6 multicast addressing for node-initiated discovery, a feature called MLD Snooping can prevent Group Manager-enabled devices from discovering each other if not configured properly. MLD Snooping is a common ether switch feature intended to reduce the amount of unnecessary IPv6 multicast traffic on a network.

If MLD Snooping is active in any network, ensure there is an MLD querier enabled so that the ether switches are kept up to date with the active Group Manager devices on the network. Alternatively, if MLD Snooping is not needed, it can be disabled. Note that some network switches have MLD Snooping enabled by default. And it is same for switching modules in the MX7000 chassis.

(i) NOTE:

For example

• To disable MLD snooping on a VLAN on a MX5108n IOM:

MX5108N-B1# configure terminal

MX5108N-B1(config)# interface vlan 194

MX5108N-B1(conf-if-vl-194)#no ipv6 mld snooping

• To enable an MLD querier in a VLAN on the MX5108n IOM:

MX5108N-B1# configure terminal

MX5108N-B1(config)# interface vlan 194

MX5108N-B1(conf-if-vl-194)#ipv6 mld snooping querier

Manage Logins

Use this section to Add New User, Change User Password and Delete User from the Group.



Group jobs including Manage Logins are one time configurations of the servers. Group manager uses SCP and jobs to make any changes. Every iDRAC in the group owns an individual job in its job queue for each Group Manager job. Group Manager does not detect changes on member iDRACs or lock member configurations.

(i) NOTE: Group jobs does not configure or override the lockdown mode for any specific iDRAC.

Leaving a group does not change local user or change settings on a member iDRAC.

Add a New User

Use this section to create and add a new user profile on all the servers in that group. A group job would be created to add the user to all servers in that group. The status of group job can be found at **GroupManager** > **Jobs** page.

(i) **NOTE:** By default iDRAC is configured with a local administrator account. You can access further information for each parameter with local administrator account.

For more information see, Configuring user accounts and privileges.

Table 36. New User Options

Option	Description
New User Information	Allows you to provide the new user's information details.
iDRAC Permissions	Allows you to define the user's role for future usage.
Advanced User Settings	Allows you to set (IPMI) user privileges and helps you to enable SNMP.

NOTE: Any member iDRAC with system lockdown enabled, that is part of the same group returns an error that the user password was not updated.

Change User Password

Use this section to change the password information for the user. You can see the **Users** detail with the **User Name**, **Role** and **Domain** information for individual user. A group job would be created to change the user password on all the servers in that group. The status of group job can be found at **GroupManager** > **Jobs** page.

If user already exists then the password can be updated. Any member iDRAC with system lockdown enabled, that is part of the group returns an error that the user password was not updated. If the user does not exist, then an error is returned to group manager indicating that the user does not exist on the system. The list of users shown in Group Manager GUI is based on the current user list on the iDRAC that is acting as the primary controller. It does not show all users for all iDRACs.

Delete User

Use this section to delete users from all the group servers. A group job would be created to delete users from all the group servers. The status of group job can be found at **GroupManager** > **Jobs** page.

If user already exists on a member iDRAC then the user can be deleted. Any member iDRAC with system lockdown enabled that is part of the group returns an error that the user is not deleted. If the user does not exist then it shows a successful deletion for that iDRAC. The list of users shown in Group Manager GUI is based on the current user list on the iDRAC which is acting as the primary controller. It does not show all users for all iDRACs.

Configure Alerts

Use this section to configure e-mail alerts. By default alerting is disabled. However, you can enable the alerting anytime. A group job would be created to apply the e-mail alert configuration to all the group servers. The status of group job can be monitored at **GroupManager** > **Jobs** page. Group manager email alert configures email alerts on all members. It sets the SMTP server settings on all members in the same group. Each iDRAC is configured separately. Email configuration is not globally saved. Current values are based on the iDRAC that is acting as a primary controller. Leaving a group does not reconfigure email alerts.

For more information on Configuring Alerts, see Configuring iDRAC to send alerts.



Table 37. Configuring alerts options

Option	Description
SMTP (Email) Server Address Settings	Allows you to configure Server IP Address, SMTP Port Number and enable the authentication. In case you are enabling authentication, you need to provide username and password.
Email Addresses	Allows you to configure multiple Email IDs to receive email notifications about system status change. You can send one test email to the configured account from the system.
Alert Categories	Allows you to select multiple alert categories to receive email notifications.

NOTE: Any member iDRAC with system lockdown enabled, that is part of the same group returns an error that the user password was not updated.

Export

Use this section to export the Group Summary to the local system. The information can be exported to a csv file format. It contains data related to each individual system in the group. Export includes the following information in csv format. Server details:

- Health
- Host Name
- iDRAC IPV4 Address
- iDRAC IPV6 Address
- Asset Tag
- Model
- iDRAC Firmware Version
- Last Status Update
- Express Service Code
- iDRAC Connectivity
- Power State
- Operating System
- Service Tag
- Node ID
- iDRAC DNS Name
- BIOS Version
- CPU Details
- System Memory(MB)
- Location Details

NOTE: In case, you are using Internet Explorer, disable the Enhanced Security settings to successfully download the csv file.

Discovered Servers View

After creating the local group, iDRAC group manager notifies all other iDRACs on the local network that a new group has been created. For iDRACs to be displayed under Discovered Servers, group manager feature should be enabled in each iDRAC. Discovered Servers View displays the list of the iDRACs detected on the same network, which can be part of any group. If an iDRAC does not show up in the discovered systems list then the user must logon to the specific iDRAC and join the group. The iDRAC that created the group will be shown as the only member in the essentials view until more iDRACs are joined to the group.

NOTE: Discovered servers view at the Group manager console allows you to onboard one or more servers listed in the view in to that group. The progress of the activity can be tracked from **GroupManager** > **Jobs**. Alternatively you can logon to



the iDRAC and select the group you would like to onboard from the drop down list to join that group. You can access the GroupManager welcome screen from iDRAC index page.

Table 38. Group onboard options

Option	Description
Onboard and Change Login	Select a specific row and select the Onboard and Change Login option to get the newly discovered systems to the group. You must provide the admin logon credentials for the new systems to join the group. If the system has the default password, you need to change it while onboarding it to a group. Group onboarding allows you to apply the same group alert settings to the new systems.
Ignore	Allows you to ignore the systems from the discovered servers list, in case you do not want to add them in any group.
Un-Ignore	Allows you to select the systems that you would like to reinstate in the discovered servers list.
Rescan	Allows you to scan and generate the list of discovered servers at any time.

Jobs View

Jobs view allows the user to track the progress of a group job, helps with simple recovery steps to correct connectivity induced failures. It also shows the history of the last group actions that were performed as an audit log. The user can use the jobs view to track the progress of the action across the group or to cancel an action that is schedule to occur in the future. The Jobs view allows the user to view the status of the last 50 jobs that have been run and any success or failures that has occurred.

Table 39. Jobs View

Option	Description
Status	Shows the job's status and the state of the ongoing job.
Job	Displays the Job's name.
ID	Displays the Job's ID.
Start Time	Displays the start time.
End Time	Displays the end time.
Actions	 Cancel — A scheduled job can be cancelled, before it moves to running state. A running job can be stopped by using the stop button. Rerun — Allows the user to rerun the job in case the job is in failure state. Remove — Allows the user to remove the completed old jobs.
Export	You can export the group job information to the local system for future references. The jobs list can be exported to csv file format. It contains data related to individual job.

NOTE: For each job entry, the list of systems provide details up to 100 systems. Each system entry contains Hostname, Service Tag, Member Job Status and Message if in case the job failed.

All Group actions that create jobs are performed on all the group members with immediate effect. You can perform the following tasks:

- Add/Edit/Remove users
- Configure email alerts



• Change group passcode and name

NOTE: Group jobs complete quickly as long as all members are online and accessible. It may take 10 minutes from job start to job complete. A job will wait and retry for up to 10 hours for the systems that are not accessible.

() NOTE: While an onboarding job is running no other Job can be scheduled. Jobs include:

- Add New User
- Change User Password
- Delete User
- Configure Alerts
- Onboard additional systems
- Change Group Passcode
- Change Group Name

Attempting to invoke another Job while an Onboarding task is active, consequences GMGR0039 error code. Once the onboarding task has made its first attempt to onboard all the new systems, Jobs can be created at any point in time.

Jobs Export

You can export the log to the local system for further references. The jobs list can be exported to a csv file format. It contains all the data related to each job.

(i) NOTE: Exported CSV files are available only in English.

Group Information Panel

Group Information panel in the top right of group manager summary view shows a consolidated group summary. Current group configuration can be edited from the Group Settings page accessible by clicking Group Settings button. It shows how many systems are there in the group. It also provides the information about the Primary and the Secondary controller of the Group.

Group Settings

Group settings page provides a listing of selected group attributes.

Table 40. Group setting attributes

Group Attribute	Description
Group Name	Displays the name of that Group.
Number of Systems	Displays the total number of systems in that group.
Created on	Displays the time stamp details.
Created by	Displays the details of Group admin.
Controlling System	Displays the Service Tag of the system, that acts as the controlling system and coordinates the group management tasks.
Backup System	Displays the Service Tag of the system, that acts as the backup system. In case the controlling system is unavailable, it takes up roles of the controlling system.

Allows the user to perform actions listed on the table below on the group. A group configuration job would be created for these actions (change group name, change group passcode, remove the members and delete the group). The status of group job can be viewed or modified from **GroupManager** > **Jobs** page.



Table 41. Group setting actions

Actions	Description
Change Name	Allows you to change the Current Group Name with a New Group Name .
Change Passcode	Allows you to change the existing group password by entering a New Group Passcode and validating that password by Reenter New Group Passcode .
Remove Systems	Allows you to remove multiple systems from the group at a time.
Delete Group	Allows you to delete the group. To use any feature of group manager, the user should have administrator privileges. Any pending jobs will be stopped in case the group is deleted.

Actions on a selected Server

On the Summary page, you can double click on a row to launch iDRAC for that server through a single sign on redirect. Ensure to turn off the popup blocker in the browser settings. You can perform following actions on the selected server by clicking appropriate item from the **More Actions** drop down list.

Table 42. Actions on a selected Server

Option	Description
Graceful Shutdown	Shuts down the operating system and powers off the system.
Cold Reboot	Powers off, then reboots the system.
Virtual Console	Launches Virtual Console with single sign on a new browser window. (i) NOTE: Disable Popup blocker from the browser to use this functionality.

Group Manager Single Sign On

All iDRACs in the group trust each other based on the shared passcode secret and shared group name. As a result an administrator user at a group member IDRAC is grant administrator level privileges at any group member iDRAC when accessed through Group Manager web interface single sign on. iDRACs logs <user>-<SVCTAG> as the user that logged on into peer members. <SVCTAG> is the service tag of the iDRAC where the user first logged in.

Group Manager Concepts — Controlling System

- Automatically selected by default the first iDRAC configured for Group Manager.
- Provides Group Manager GUI workflow.
- Keeps track of all members.
- Coordinates tasks.
- If a user logs in to any member and clicks on Open Group Manager the browser will be redirected to the primary controller.

Group Manager Concepts — Backup System

- Primary controller automatically selects a secondary controller to take over if the primary goes offline for an extended period of time (10 mins or more).
- If both primary and secondary goes offline for an extended duration (for more than 14 mins) a new primary and secondary controller gets elected.
- Keeps a copy of the group manager cache of all the groups members and tasks.
- The controlling system and backup system are automatically determined by group manager.



• No user configuration or involvement required.

iDRAC Group Firmware Update

For iDRAC group firmware update, from the DUP file from a local directory, perform the following steps:

- 1. Access group manager console essential view and click Update iDRAC Firmware under summary view.
- 2. From the firmware update dialog box displayed, browse and select the local iDRAC DUP file to be installed. Click Upload.
- **3.** File is uploaded to iDRAC and verified for integrity.
- 4. Confirm the firmware update. Group iDRAC firmware update job is scheduled for immediate execution. If Group Manager has other group jobs running, then it is executed after the previous job is completed.
- 5. You can track iDRAC update job execution from group jobs view.

(i) NOTE: This feature is only supported on iDRAC version 3.50.50.50 and above.



Managing logs

iDRAC provides Lifecycle log that contains events related to system, storage devices, network devices, firmware updates, configuration changes, license messages, and so on. However, the system events are also available as a separate log called System Event Log (SEL). The lifecycle log is accessible through iDRAC Web interface, RACADM, and WSMan interface.

When the size of the lifecycle log reaches 800 KB, the logs are compressed and archived. You can only view the non-archived log entries, and apply filters and comments to non-archived logs. To view the archived logs, you must export the entire lifecycle log to a location on your system.

Topics:

- Viewing System Event Log
- Viewing Lifecycle log
- Exporting Lifecycle Controller logs
- Adding work notes
- Configuring remote system logging

Viewing System Event Log

When a system event occurs on a managed system, it is recorded in the System Event Log (SEL). The same SEL entry is also available in the LC log.

(i) NOTE: SEL and LC logs may have mismatch in timestamp when iDRAC is rebooting.

Viewing System Event Log using web interface

To view the SEL, in iDRAC Web interface, go to Maintenance > System Event Log.

The **System Event Log** page displays a system health indicator, a time stamp, and a description for each event logged. For more information, see the *iDRAC Online Help*.

Click Save As to save the SEL to a location of your choice.

NOTE: If you are using Internet Explorer and if there is a problem when saving, download the Cumulative Security Update for Internet Explorer. You can download it from the Microsoft Support website at **support.microsoft.com**.

To clear the logs, click Clear Log.

(i) NOTE: Clear Log only appears if you have Clear Logs permission.

After the SEL is cleared, an entry is logged in the Lifecycle Controller log. The log entry includes the user name and the IP address from where the SEL was cleared.

Viewing System Event Log using RACADM

To view the SEL: racadm getsel <options> If no arguments are specified, the entire log is displayed. To display the number of SEL entries: racadm getsel -i To clear the SEL entries: racadm clrsel For more information, see *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Managing logs 193



Viewing System Event Log using iDRAC settings utility

You can view the total number of records in the System Event Log (SEL) using the iDRAC Settings Utility and clear the logs. To do this:

- In the iDRAC Settings Utility, go to System Event Log. The iDRAC Settings.System Event Log displays the Total Number of Records.
- 2. To clear the records, select Yes. Else, select No.
- 3. To view the system events, click **Display System Event Log**.
- 4. Click Back, click Finish, and then click Yes.

Viewing Lifecycle log

Lifecycle Controller logs provide the history of changes related to components installed on a managed system. You can also add work notes to each log entry.

The following events and activities are logged:

- All
- System Health System Health category represents all the alerts that are related to hardware within the system chassis.
- Storage Storage Health category represents alerts that are related to the storage subsystem.
- Updates Update category represents alerts that are generated due to firmware/Driver upgrades/downgrades.
- Audit Audit category represents the audit log.
- Configuration Configuration category represents alerts that are related to hardware, firmware and software configuration changes.
- Work Notes

When you log in to or log out of iDRAC using any of the following interfaces, the log-in, log-out, or login failure events are recorded in the Lifecycle logs:

- SSH
- Web interface
- RACADM
- Redfish
- IPMI over LAN
- Serial
- Virtual console
- Virtual media

You can view and filter logs based on the category and severity level. You can also export and add a work note to a log event.

(i) NOTE: Lifecycle logs for Personality Mode change is generated only during the warm boot of the host.

If you initiate configuration jobs using RACADM CLI or iDRAC web interface, the Lifecycle log contains information about the user, interface used, and the IP address of the system from which you initiate the job.

NOTE: On MX platform, Lifecycle Controller logs multiple job IDs for configuration or installation jobs created using OME - Modular. For more information on the jobs performed, see the OME - Modular logs.

Viewing Lifecycle log using web interface

To view the Lifecycle Logs, click **Maintenance** > Lifecycle Log. The Lifecycle Log page is displayed. For more information about the options, see the *iDRAC Online Help*.

Filtering Lifecycle logs

You can filter logs based on category, severity, keyword, or date range.

To filter the lifecycle logs:

1. In the Lifecycle Log page, under the Log Filter section, do any or all of the following:



- Select the Log Type from the drop-down list.
- Select the severity level from the **Severity** drop-down list.
- Enter a keyword.
- Specify the date range.
- 2. Click Apply.

The filtered log entries are displayed in Log Results.

Adding comments to Lifecycle logs

To add comments to the Lifecycle logs:

- In the Lifecycle Log page, click the + icon for the required log entry. The Message ID details are displayed.
- Enter the comments for the log entry in the Comment box. The comments are displayed in the Comment box.

Viewing Lifecycle log using RACADM

To view Lifecycle logs, use the lclog command.

For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Exporting Lifecycle Controller logs

You can export the entire Lifecycle Controller log (active and archived entries) in a single zipped XML file to a network share or to the local system. The zipped XML file extension is .xml.gz. The file entries are ordered sequentially based on their sequence numbers, ordered from the lowest sequence number to the highest.

Exporting Lifecycle Controller logs using web interface

To export the Lifecycle Controller logs using the Web interface:

- 1. In the Lifecycle Log page, click Export.
- 2. Select any of the following options:
 - Network Export the Lifecycle Controller logs to a shared location on the network.
 - Local Export the Lifecycle Controller logs to a location on the local system.
 - **NOTE:** While specifying the network share settings, it is recommended to avoid special characters for user name and password or percent encode the special characters.

For information about the fields, see the iDRAC Online Help.

3. Click Export to export the log to the specified location.

Exporting Lifecycle Controller logs using RACADM

To export the Lifecycle Controller logs, use the lclog export command.

For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Adding work notes

Each user who logs in to iDRAC can add work notes and this is stored in the lifecycle log as an event. You must have iDRAC logs privilege to add work notes. A maximum of 255 characters are supported for each new work note.

(i) NOTE: You cannot delete a work note.

To add a work note:



- In the iDRAC Web interface, go to Dashboard > Notes > add note. The Work Notes page is displayed.
- 2. Under Work Notes, enter the text in the blank text box.

(i) NOTE: It is recommended not to use too many special characters.

3. Click Save.

The work note is added to the log. For more information, see the *iDRAC Online Help*.

Configuring remote system logging

You can send lifecycle logs to a remote system. Before doing this, make sure that:

- There is network connectivity between iDRAC and the remote system.
- The remote system and iDRAC is on the same network.

Configuring remote system logging using web interface

To configure the remote syslog server settings:

- In the iDRAC Web interface, go to Configuration > System Settings > Remote Syslog Settings. The Remote Syslog Settings page is displayed
- 2. Enable remote syslog, specify the server address, and the port number. For information about the options, see the *iDRAC Online Help*.
- 3. Click Apply.

The settings are saved. All logs written to the lifecycle log are also simultaneously written to configured remote server(s).

Configuring remote system logging using RACADM

To configure the remote system-logging settings, use the set command with the objects in the iDRAC.SysLog group.

For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.



Monitoring and managing power in iDRAC

You can use iDRAC to monitor and manage the power requirements of the managed system. This helps to protect the system from power outages by appropriately distributing and regulating the power consumption on the system.

The key features are:

- **Power Monitoring** View the power status, history of power measurements, the current averages, peaks, and so on for the managed system.
- **Power Capping** View and set the power cap for the managed system, including displaying the minimum and maximum potential power consumption. This is a licensed feature.
- **Power Control** Enables you to remotely perform power control operations (such as, power on, power off, system reset, power cycle, and graceful shutdown) on the managed system.
- Power Supply Options Configure the power supply options such as redundancy policy, hot spare, and power factor correction.

Topics:

- Monitoring power
- Setting warning threshold for power consumption
- Executing power control operations
- Power capping
- Configuring power supply options
- Enabling or disabling power button
- Multi-Vector Cooling

Monitoring power

iDRAC monitors the power consumption in the system continuously and displays the following power values:

- Power consumption warning and critical thresholds.
- Cumulative power, peak power, and peak amperage values.
- Power consumption over the last hour, last day or last week.
- Average, minimum, and maximum power consumption.
- Historical peak values and peak timestamps.
- Peak headroom and instantaneous headroom values (for rack and tower servers).

NOTE: The histogram for the system power consumption trend (hourly, daily, weekly) is maintained only while iDRAC is running. If iDRAC is restarted, the existing power consumption data is lost and the histogram is restarted.

(i) NOTE: After iDRAC firmware update or reset, the power consumption graph will be wiped / reset.

Monitoring performance index of CPU, memory, and input output modules using web interface

To monitor the performance index of CPU, memory, and I/O modules, in the iDRAC web interface, go to **System** > **Performance**.

- System Performance section Displays the current reading and the warning reading for CPU, Memory and I/O utilization index, and system level CUPS index in a graphical view.
- System Performance Historical Data section:
 - Provides the statistics for CPU, memory, IO utilization, and the system level CUPS index. If the host system is powered off, then the graph displays the power off line below 0 percent.

Monitoring and managing power in iDRAC 197

Inserido ao protocolo **22.951.206-4** por: **Pedro Henrique Golin Linhares** em: 31/03/2025 12:49. A autenticidade deste documento pode ser validada no endereço: https://www.eprotocolo.pr.gov.br/spiweb/validarDocumento com o código: 692a5ef82beb173ed33dad4ee6505546.



- You can reset the peak utilization for a particular sensor. Click **Reset Historical Peak**. You must have Configure privilege to reset the peak value.
- Performance Metrics section:
 - Displays status and present reading
 - Displays or specifies the warning threshold utilization limit. You must have server configure privilege to set the threshold values.

For information about the displayed properties, see the *iDRAC Online Help*.

Monitoring performance index for of CPU, memory, and input output modules using RACADM

Use the **SystemPerfStatistics** sub command to monitor performance index for CPU, memory, and I/O modules. For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Setting warning threshold for power consumption

You can set the warning threshold value for the power consumption sensor in the rack and tower systems. The warning/critical power threshold for rack and tower systems may change, after the system is power-cycled, based on PSU capacity and redundancy policy. However, the warning threshold must not exceed the critical threshold even if Power Supply Unit capacity of the redundancy policy is changed.

The warning power threshold for blade systems is set to power allocation of CMC(for non-MX platforms) or OME Modular (for MX platforms).

If reset to default action is performed, the power thresholds will be set to default.

You must have Configure user privilege to set the warning threshold value for power consumption sensor.

(i) NOTE: The Warning Threshold value is reset to the default value after performing a racreset or an iDRAC update.

Setting warning threshold for power consumption using web interface

- 1. In the iDRAC Web interface, go to System > Overview > Present Power Reading and Thresholds.
- 2. In the **Present Power Reading and Thresholds** section, click **Edit Warning Threshold**. The **Edit Warning Threshold** page is displayed.
- 3. In the Warning Threshold column, enter the value in Watts or BTU/hr. The values must be lower than the Failure Threshold values. The values are rounded off to the nearest value that is divisible by 14. If you enter Watts, the system automatically calculates and displays the BTU/hr value. Similarly, if you enter BTU/hr, the value for Watts is displayed.
- 4. Click Save. The values are configured.

Executing power control operations

iDRAC enables you to remotely perform a power-on, power off, reset, graceful shutdown, Non-Masking Interrupt (NMI), or power cycle using the Web interface or RACADM.

You can also perform these operations using Lifecycle Controller Remote Services or WSMan. For more information, see the Lifecycle Controller Remote Services Quick Start Guide available at https://www.dell.com/idracmanuals and the Dell Power State Management Profile document available at https://www.dell.com/support.

Server power-control operations initiated from iDRAC are independent of the power-button behavior configured in the BIOS. You can use the PushPowerButton function to gracefully shut down the system, or power it on, even if the BIOS is configured to do nothing when the physical power button is pressed.

198 Monitoring and managing power in iDRAC



Executing power control operations using web interface

To perform power control operations:

- 1. In iDRAC web interface, go to **Configuration** > **Power Management** > **Power Control**. The **Power Control** options are displayed.
- **2.** Select the required power operation:
 - Power On System
 - Power Off System
 - NMI (Non-Masking Interrupt)
 - Graceful Shutdown
 - Reset System (warm boot)
 - Power Cycle System (cold boot)
- 3. Click Apply. For more information, see the *iDRAC Online Help*.

Executing power control operations using RACADM

To perform power actions, use the **serveraction** command.

For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Power capping

You can view the power threshold limits that covers the range of AC and DC power consumption that a system under heavy workload presents to the datacenter. This is a licensed feature.

Power capping in Blade servers

Before a blade server turns on, based on limited hardware inventory, iDRAC provides the power requirements of the blade server to the chassis manager. If the power consumption increases over time and if the server consumes power to its maximum allocation, iDRAC requests CMC (for non-MX platforms) or OME Modular (for MX platforms) to increase the maximum potential power. This results in an increase in the power delivery, however, the power delivery does not reduce if the consumption decreases.

After the system is powered on and initialized, iDRAC calculates a new power requirement based on the actual hardware configuration. The system stays powered on even if the CMC (not for MX platforms) or OME Modular (not for MX platforms) fails to allocate new power request.

CMC or OME Modular reclaims any unused power from lower priority servers and allocates that power to a higher-priority infrastructure module or a server.

Viewing and configuring power cap policy

When power cap policy is enabled, it enforces a user-defined power limits on the system. If power-capping is not enabled, the default hardware power-protection policy is used. This power-protection policy is independent of the user-defined policy. The system performance is dynamically adjusted to maintain power consumption close to the specified threshold.

Actual power consumption depends on the workload. It may momentarily exceed the threshold until performance adjustments are completed. For example, consider a system that has a minimum and maximum Potential Power Consumption values of 500 W and 700 W respectively. You can specify a Power Budget Threshold to reduce consumption to 525 W. When this power budget is configured, the performance of the system is dynamically adjusted to maintain power consumption of 525 W or less.

If you set a very low power cap or if the ambient temperature is unusually high, power consumption may temporarily exceed the power-cap while the system is powering up or being reset.

If the power cap value is set lower than the minimum recommended threshold, iDRAC may not be able maintain the requested power cap.

You can specify the value in Watts, BTU/hr, or as a percentage of the recommended maximum power limit.



When setting the power cap threshold in BTU/hr, the conversion to Watts is rounded off to the nearest integer. When the power cap threshold are read from the system, the Watts to BTU/hr conversion is also rounded off. Because of the rounding off, the actual values may slightly differ.

Configuring power cap policy using web interface

To view and configure the power policies:

- 1. In iDRAC Web interface, go to **Configuration** > **Power Management** > **Power Cap Policy**. The current power policy limit is displayed under the **Power Cap Limits** section.
- 2. Select Enable under Power Cap.
- 3. Under **Power Cap Limits** section, enter the power limit within recommended range in Watts and BTU/hr or the maximum % of recommended system limit.
- 4. Click Apply to apply the values.

Configuring power cap policy using RACADM

To view and configure the current power cap values, use the following objects with the set command:

- System.Power.Cap.Enable
- System.Power.Cap.Watts
- System.Power.Cap.Btuhr
- System.Power.Cap.Percent

For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Configuring power cap policy using iDRAC settings utility

To view and configure power policies:

1. In iDRAC Settings utility, go to Power Configuration.

(i) NOTE: The Power Configuration link is available only if the server power supply unit supports power monitoring.

The iDRAC Settings Power Configuration page is displayed.

- 2. Select Enabled to enable the Power Cap Policy Else, select Disabled.
- **3.** Use the recommended settings, or under **User Defined Power Cap Policy**, enter the required limits. For more information about the options, see the *iDRAC Settings Utility Online Help*.
- Click Back, click Finish, and then click Yes. The power cap values are configured.

Configuring power supply options

You can configure the power supply options such as redundancy policy, hot spare, and power factor correction.

Hot spare is a power supply feature that configures redundant Power Supply Units (PSUs) to turn off depending on the server load. This allows the remaining PSUs to operate at a higher load and efficiency. This requires PSUs that support this feature, so that it quickly powers ON when needed.

In a two PSU system, either PSU1 or PSU2 can be configured as the primary PSU.

After Hot Spare is enabled, PSUs can become active or go to sleep based on load. If Hot Spare is enabled, asymmetric electrical current sharing between the two PSUs is enabled. One PSU is *awake* and provides the majority of the current; the other PSU is in sleep mode and provides a small amount of the current. This is often called 1 + 0 with two PSUs and hot spare enabled. If all PSU-1s are on Circuit-A and all PSU-2s are on Circuit-B, then with hot spare enabled (default hot spare factory configuration), Circuit-B has much less load and triggers the warnings. If hot spare is disabled, the electrical current sharing is 50-50 between the two PSUs, the Circuit-A and Circuit-B normally has the same load.

Power factor is the ratio of real power consumed to the apparent power. When power factor correction is enabled, the server consumes a small amount of power when the host is OFF. By default, power factor correction is enabled when the server is shipped from the factory.

200 Monitoring and managing power in iDRAC



Configuring power supply options using web interface

To configure the power supply options:

- 1. In iDRAC Web interface, go to Configuration > Power Management > Power Configuration.
- 2. Under **Power Redundancy Policy**, select the required options. For more information, see *iDRAC Online Help*.
- 3. Click Apply. The power supply options are configured.

Configuring power supply options using RACADM

To configure the power supply options, use the following objects with the get/set command:

- System.Power.RedundancyPolicy
- System.Power.Hotspare.Enable
- System.Power.Hotspare.PrimaryPSU
- System.Power.PFC.Enable

For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Configuring power supply options using iDRAC settings utility

To configure the power supply options:

1. In iDRAC Settings utility, go to **Power Configuration**.

(i) NOTE: The Power Configuration link is available only if the server power supply unit supports power monitoring.

The iDRAC Settings Power Configuration page is displayed.

- 2. Under Power Supply Options:
 - Enable or disable power supply redundancy.
 - Enable or disable hot spare.
 - Set the primary power supply unit.
 - Enable or disable power factor correction. For more information about the options, see the *iDRAC Settings Utility Online Help*.
- Click Back, click Finish, and then click Yes. The power supply options are configured.

Enabling or disabling power button

To enable or disable the power button on the managed system:

- 1. In iDRAC Settings utility, go to **Front Panel Security**.
- The **iDRAC Settings Front Panel Security** page is displayed.
- $\label{eq:constraint} \textbf{2. Select Enabled} \text{ to enable the power button or } \textbf{Disabled} \text{ to disable it.}$
- **3.** Click **Back**, click **Finish**, and then click **Yes**. The settings are saved.

Multi-Vector Cooling

Multi-Vector Cooling implements multi-prong approach to Thermal Controls in Dell EMC Server Platforms. You can configure multi-vector cooling options through iDRAC web interface by navigating to **Configuration** > **System Settings** > **Hardware Settings** > **Fan Configuration**. It includes (but not limited to):

- Large set of sensors (thermal, power, inventory etc.) that allows accurate interpretation of real-time system thermal state at various locations within the server. It displays only a small subset of sensors that are relevant to users need based on the configuration.
- Intelligent and adaptive closed loop control algorithm optimizes fan response to maintain component temperatures. It also conserves fan power, airflow consumption, and acoustics.

Monitoring and managing power in iDRAC 201



- Using fan zone mapping, cooling can be initiated for the components when it requires. Thus, it results maximum performance without compromising the efficiency of power utilization.
- Accurate representation of slot by slot PCIe airflow in terms of LFM metric (Linear Feet per Minute an accepted industry standard on how PCIe card airflow requirement is specified). Display of this metric in various iDRAC interfaces allows user to:
 - 1. know the maximum LFM capability of each slot within the server.
 - 2. know what approach is being taken for PCIe cooling for each slot (airflow controlled, temperature controlled).
 - 3. know the minimum LFM being delivered to a slot, if the card is a 3rd Party Card (user defined custom card).
 - **4.** dial in custom minimum LFM value for the 3rd Party Card allowing more accurate definition of the card cooling needs for which the user is better aware of through their custom card specification.
- Displays real-time system airflow metric (CFM, cubic feet per minute) in various iDRAC interfaces to the user to enable datacenter airflow balancing based on aggregation of per server CFM consumption.
- Allows custom thermal settings like Thermal Profiles (Maximum Performance vs. Maximum Performance per Watt, Sound
- Cap), custom fan speed options (minimum fan speed, fan speed offsets) and custom Exhaust Temperature settings.
 - 1. Most of these settings allow additional cooling over the baseline cooling generated by thermal algorithms and do not allow fan speeds to go below system cooling requirements.
 - () NOTE: One exception to above statement is for fan speeds that are added for 3rd Party PCIe cards. The thermal algorithm provision airflow for 3rd party cards may be more or less than the actual card cooling needs and customer may fine tune the response for the card by entering the LFM corresponding to the 3rd Party Card.
- Custom Exhaust Temperature option limits exhaust temperature to customer desired settings.
 NOTE: It is important to note that with certain configurations and workloads, it may not be physically possible to reduce exhaust below a desired set point (e.g. Custom exhaust setting of 45C with a high inlet temp {e.g. 30C} and
 - a loaded config {high system power consumption, low airflow}).
- **3.** Sound Cap option is new in the 14th generation of PowerEdge server. It limits CPU power consumption and controls fan speed and acoustical ceiling. This is unique for acoustical deployments and may result in reduced system performance.
- System layout and design enables increased airflow capability (by allowing high power) and dense system configurations. It
 provides less system restrictions and increased feature density.
 - 1. Streamlined airflow permits efficient airflow to fan power consumption ratio.
- Custom fans are designed for higher efficiency, better performance, longer life and less vibration. It also delivers better acoustics outcome.
 - **1.** Fans are capable of long life (in general it may run for more than 5 years), even if it runs at full speed all the time.
- Custom heat-sinks are designed for optimize component cooling at minimum (required) airflow yet supports high performance CPUs.



iDRAC Direct Updates

iDRAC provides out of band ability to update the firmware of various components of a PowerEdge server. iDRAC direct update helps in eliminating staged jobs during updates.

iDRAC used to have staged updates to initiate firmware update of the components. From this release, Direct updates have been applied to PSU and Backplane. With the use of Direct Updates and Backplane can have quicker updates. In case of PSU, one reboot (for initializing the updates) is avoided and the update can happen in single reboot.

With Direct update feature in iDRAC, you can eliminate the first reboot to initiate the updates. The second reboot will be controlled by the device itself and iDRAC notifies the user if there is need for a separate reset via job status.

iDRAC Direct Updates 203



Inventorying, monitoring, and configuring network devices

You can inventory, monitor, and configure the following network devices:

- Network Interface Cards (NICs)
- Converged Network Adapters (CNAs)
- LAN On Motherboards (LOMs)
- Network Daughter Cards (NDCs)
- Mezzanine cards (only for blade servers)

Before you disable NPAR or an individual partition on CNA devices, ensure that you clear all I/O identity attributes (Example: IP address, virtual addresses, initiator, and storage targets) and partition-level attributes (Example: Bandwidth allocation). You can disable a partition either by changing the VirtualizationMode attribute setting to NPAR or by disabling all personalities on a partition.

Depending on the type of installed CNA device, the settings of partition attributes may not be retained from the last time the partition was active. Set all I/O identity attributes and partition-related attributes when enabling a partition. You can enable a partition by either changing the VirtualizationMode attribute setting to NPAR or by enabling a personality (Example: NicMode) on the partition.

Topics:

- Inventorying and monitoring network devices
- Inventorying and monitoring FC HBA devices
- Inventorying and monitoring SFP Transceiver devices
- Telemetry Streaming
- Serial Data Capture
- Dynamic configuration of virtual addresses, initiator, and storage target settings

Inventorying and monitoring network devices

You can remotely monitor the health and view the inventory of the network devices in the managed system.

For each device, you can view the following information of the ports and enabled partitions:

- Link Status
- Properties
- Settings and Capabilities
- Receive and Transmit Statistics
- iSCSI, FCoE initiator, and target information

Monitoring network devices using web interface

To view the network device information using Web interface, go to **System** > **Overview** > **Network Devices**. The **Network Devices** page is displayed. For more information about the displayed properties, see *iDRAC Online Help*.

Monitoring network devices using RACADM

To view information about network devices, use the hwinventory and nicstatistics commands.

For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Additional properties may be displayed when using RACADM or WSMan in addition to the properties displayed in the iDRAC web interface.

204 Inventorying, monitoring, and configuring network devices



Connection View

Manually checking and troubleshooting the servers' networking connections is unmanageable in a datacenter environment. iDRAC9 streamlines the job with iDRAC Connection View. This feature allows you to remotely check and troubleshoot network connections from the same centralized GUI that you are using for deploying, updating, monitoring, and maintaining the servers. Connection View in iDRAC9 provide details of the physical mapping of switch ports to server's network ports and iDRAC (integrated Dell Remote Access Controller) dedicated port connections. All supported network cards are visible in Connection View, irrespective of the brand.

Instead of manually checking and troubleshooting the server's networking connections, you can view and manage network cable connections remotely.

The Connection View provides the information of the switch ports which are connected to the server ports, and iDRAC dedicated port. The server network ports include those on PowerEdge LOM, NDC, Mezz cards, and PCIe add-in cards.

To View network devices connection view, navigate to **System** > **Overview** > **Network Device** > **Connection View** to view the Connection View.

Also, you can click **iDRAC Settings** > **Connectivity** > **Network** > **Common Settings** > **Connection View** to enable or disable the connection view.

Connection View can be explored with racadm SwitchConnection View command and it can also be viewed with command.

Field or Option	Description
Enabled	Select Enabled , to enable Connection View. By default the Enabled option is selected.
State	Displays Enabled , if you enable the connection view option from the Connection View from iDRAC settings.
Switch Connection ID	Displays the LLDP chassis ID of the switch through which the device port is connected.
Switch Port Connection ID	Displays the LLDP port ID of the switch port to which the device port is connected.

() NOTE: Switch Connection ID and Switch Port Connection ID are available once the Connection View is enabled and the Link is connected. The associated network card needs to be compatible with the Connection View. Only users with iDRAC Configure privilege can modify the Connection View settings.

From iDRAC9 4.00.00.00 and later versions, iDRAC supports to send standard LLDP packets to external switches. This provides options to discover iDRACs on the network. iDRAC sends two types of LLDP packets to the outbound network:

• **Topology LLDP** - In this feature, the LLDP packet goes through all the supported server NIC ports so that an external switch can locate the originating server, NDC port[NIC FQDD], IOM location in the chassis, blade chassis service tag etc. From iDRAC9 4.00.00.00 and later versions, Topology LLDP is available as an option for all PowerEdge servers. The LLDP packets contain server network device connectivity information and are used by I/O modules and external switches to update their configuration.

() NOTE:

- The Topology LLDP must be enabled for the MX chassis configuration to function properly.
- The Topology LLDP is not supported on 1GbE controllers and select 10GbE controllers (Intel X520, QLogic 578xx).
- **Discovery LLDP** In this feature, the LLDP packet goes only through the active iDRAC NIC port in use (dedicated NIC or shared LOM), so an adjacent switch can locate iDRAC connection port in the switch. Discovery LLDP is specific only to the active iDRAC network port and not will not be seen in all the Network ports in the server. Discovery LLDP will have some details of the idrac like IP address, MAC address, service tag etc., so that a switch can automatically discover iDRAC devices connected to it and some data of iDRAC.

(i) NOTE: If Virtual MAC Address is cleared on a port/partition, then Virtual MAC Address would be same as MAC Address.

To enable or disable the Topology LLDP, navigate to **iDRAC Settings** > **Connectivity** > **Network** > **Common Settings** > **Topology LLDP** to enable or disable the Topology LLDP. By default, it is enabled for MX servers and disabled for all other servers.

To enable or disable the iDrac Discovery LLDP, navigate to **iDRAC Settings** > **Connectivity** > **Network** > **Common Settings** > **iDrac Discovery LLDP**. By default, the Enabled option is selected.

LLDP packet originated from idrac can be viewed from switch using the command: show lldp neighbors.



Refresh Connection View

Use Refresh Connection View to get the latest information of Switch Connection ID and Switch Port Connection ID.

() NOTE: If iDRAC has switch connection and switch port connection information for server network port or iDRAC network port and due to some reason, the switch connection and switch port connection information is not refreshed for 5min, then the switch connection and switch port connection information is shown as stale (last known good data) data for all user interfaces. In the UI, you see yellow bang which is a natural representation and it does not indicate any warning.

Connection View Possible Values

Possible Connection View Data	Description
Feature Disabled	Connection view feature is disabled, to view the connection view data enable the feature.
No Link	Indicates that the link associated with network controller port is down.
Not Available	LLDP is not enabled on the switch. Check whether LLDP is enabled on the switch port.
Not Supported	Network controller does not support Connection view feature.
Stale Data	Last known good data, either the Network controller port link is down or the system is powered off. Use the refresh option to refresh the connection view details to get the latest data.
Valid Data	Displays the Valid Switch Connection ID and the Switch Port Connection ID information.

Connection View Supported Network Controllers

Following cards or controllers support Connection View feature.

Manufacturer	Туре				
Broadcom	 57414 rNDC 25GE 57416/5720 rNDC 10GbE 57412/5720 rNDC 10GbE 57414 PCIe FH/LP 25GE 57412 PCIe FH/LP 10GbE 57416 PCIe FH/LP 10GbE 				
Intel	 X710 bNDC 10Gb X710 DP PCIe 10Gb X710 QP PCIe 10Gb X710 + I350 rNDC 10Gb+1Gb X710 rNDC 10Gb X710 bNDC 10Gb XL710 bNDC 10Gb XL710 PCIe 40Gb XL710 OCP Mezz 10Gb X710 PCIe 10Gb 				
Mellanox	 MT27710 rNDC 40Gb MT27710 PCIe 40Gb MT27700 PCIe 100Gb 				
QLogic	 QL41162 PCIe 10GE 2P QL41112 PCIe 10GE 2P QL41262 PCIe 25GE 2P 				

206 Inventorying, monitoring, and configuring network devices



Inventorying and monitoring FC HBA devices

You can remotely monitor the health and view the inventory of the Fibre Channel Host Bus Adapters (FC HBA) devices in the managed system. The Emulex and QLogic FC HBAs are supported. For each FC HBA device, you can view the following information for the ports:

- Link Status and Information
- Port Properties
- Receive and Transmit Statistics

(i) NOTE: Emulex FC8 HBAs are not supported.

Monitoring FC HBA devices using web interface

To view the FC HBA device information using Web interface, go to **System** > **Overview** > **Network Devices** > **Fibre Channel**. For more information about the displayed properties, see *iDRAC Online Help*.

The page name also displays the slot number where the FC HBA device is available and the type of FC HBA device.

Monitoring FC HBA devices using RACADM

To view the FC HBA device information using RACADM, use the hwinventory command.

For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Inventorying and monitoring SFP Transceiver devices

You can remotely monitor the health and view the inventory of SFP transceiver devices connected to the system. Following are the supported transceivers:

- SFP
- SFP+
- SFP28
- SFP-DD
- QSFP
- QSFP+
- QSFP28
- QSFP-DD
- Base-T modules
- AOC & DAC cables
- RJ-45 Base-T connected with Ethernet
- Fiber channel
- IB adapter ports

Most useful transceiver information are Serial number and Part number from transceiver EPROM. These would allow to verify the remotely installed transceivers, when troubleshooting connectivity issues. For each SFP Transceiver device, you can view the following information for the ports:

- Vendor Name
- Part Number
- Revision
- Serial Number
- Device Identifier / Type Info
- Cable length (in meter)



Monitoring SFP Transceiver devices using web interface

To view theSFP Transceiver device information using Web interface, go to **System** > **Overview** > **Network Devices** and click on particular device. For more information about the displayed properties, see *iDRAC Online Help*.

The page name also displays the slot number where the transceiver device is available under Port statistics.

Monitoring SFP Transceiver devices using RACADM

To view the SFP Transceiver device information using RACADM, use the hwinventory command.

For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Telemetry Streaming

Telemetry enables users to collect and stream real-time device metrics, events, and data logs from a PowerEdge server to a subscribed external client or server application. Using Telemetry, you can set the type and frequency of reports that needs to be generated.

(i) NOTE: The feature is supported on all the platforms and it requires iDRAC Datacenter license.

Telemetry is one-to-many solution for collecting and streaming the live system data from one or more PowerEdge servers (iDRAC) to a centralized 'Remote Server Monitoring, Analysis, and Alerting service'. The feature also supports on-demand data collection of the data.

The telemetry data includes metrics/inventory and logs/events. The data can be streamed (pushed out) or collected (pulled) from iDRAC to or by remote consumers like Redfish client and Remote Syslog Server. The telemetry data is also provided to the iDRAC SupportAssist data collector on demand. The data collection and report is based on predefined Redfish telemetry metrics, trigger, and report definitions. The telemetry streaming settings can be configured through RACADM, Redfish, and Server Configuration Profile (SCP).

To configure Telemetry, enable or select the required device reports or logs that define the behavior and frequency of data streaming. Go to **Configuration** > **System Settings** page to configure Telemetry. Data streaming is automatic until the Telemetry is disabled.

Туре	Metric Group	Inventory	Sensor	Statistics	Configuration	Metrics
I/O Devices	NICs	No	Yes	Yes	No	No
	FC HBAs	No	Yes	Yes	No	No
Server Devices	CPUs	No	Yes	No	No	Yes
	Memory	No	Yes	No	No	Yes
	Fans	No	Yes	No	No	No
	PSUs	No	No	No	No	Yes
	Sensors	No	Yes	No	No	No
Environmental	Thermal	No	Yes	No	No	Yes
	Power	No	No	Yes	No	Yes
	Performance	No	No	Yes	No	No
Accelerators	GPUs	No	No	Yes	No	Yes

Following table describes the metric reports that can be generated using telemetry:

To know about the field descriptions of Telemetry section, see *iDRAC Online Help*.

() NOTE:

• StorageDiskSMARTDATA is only supported on SSD drives with SAS/SATA bus protocol and behind the BOSS controller.


- StorageSensor data is reported only for the drives in Ready / Online / Non-RAID mode and not behind the BOSS controller.
- NVMeSMARTData is only supported for SSD (PCleSSD / NVMe Express) drives with PCle bus protocol (not behind SWRAID).
- GPGPUStatistics data is only available in specific GPGPU models that support ECC memory capability.
- PSUMetrics is not available on modular platforms.
- Fan Power and PCIe Power Metrics may be displayed as 0 for some platforms.
- CUPS report has been renamed to SystemUsage in 4.40.00.00 release and it's supported on both INTEL and AMD platforms.

Telemetry Workflow:

- 1. Install Datacenter license, if not installed already.
- 2. Configure global Telemetry settings including Enabling the telemetry and Rsyslog server network address and port using RACADM, Redfish, SCP, or iDRAC GUI.
- **3.** Configure the following Telemetry report streaming parameters on the required device report or log using either RACADM or Redfish interface:
 - EnableTelemetry
 - ReportInterval
 - ReportTriggers

(i) NOTE: Enable iDRAC Alerts and Redfish events for the specific hardware for which you need telemetry reports.

- **4.** Redfish client makes subscription request to the Redfish EventService on iDRAC.
- 5. iDRAC generates and pushes the metric report or log/event data to the subscribed client when the predefined trigger conditions are met.

Feature Constraints:

- 1. For security reasons, iDRAC supports only HTTPS-based communication to the client.
- 2. For stability reasons, iDRAC supports up to eight subscriptions.
- **3.** Deletion of subscriptions is supported through Redfish interface only, even for the manual deletion by the Admin.

Behavior of Telemetry feature:

- iDRAC generates and pushes (HTTP POST) the Metric Report or log/event data to all the subscribed clients to the destination specified in the subscription when the predefined trigger conditions are met. The clients receive new data only upon successful subscription creation.
- The metric data includes the timestamp in ISO format, UTC time (ends in 'Z'), at the time of data collection from source.
- Clients can terminate a subscription by sending an HTTP DELETE message to the URI of the subscription resource through the Redfish interface.
- If the subscription is deleted either by iDRAC or the client, then iDRAC does not send (HTTP POST) reports. If the number of delivery errors exceeds predefined thresholds, then iDRAC may delete a subscription.
- If a user has Admin privilege, they can delete the subscriptions but only through Redfish interface.
- Client is notified about the termination of a subscription by iDRAC by sending 'Subscription terminated' event as the last message.
- Subscriptions are persistent and can remain even after iDRAC restarts. But, they can be deleted either by performing racresetcfg or LCwipe operations.
- User interfaces like RACADM, Redfish, SCP, and iDRAC display the current status of the client subscriptions.

Serial Data Capture

iDRAC allows you to capture console redirection serial for later retrieval with the use of Serial Data Capture feature. This feature requires iDRAC Datacenter license.

The purpose of Serial Data Capture feature is to capture the system serial data and store it so that the customer can later retrieve it for debugging purpose.

You can enable or disable a serial data capture using RACADM, Redfish, iDRAC interfaces. When this attribute is enabled, iDRAC will capture serial traffic received on Host Serial Device2 irrespective of serial Mux mode settings.

To enable / disable Serial Data Capture using iDRAC GUI, go to **Maintainance** > **Diagnostics** > **Serial Data Logs** page, and check the box to enable or disable.



() NOTE:

- This attribute is persistent over iDRAC reboot.
- Firmware reset to default will disable this feature.
- While Serial Data capture is enabled, the buffer keeps getting appended with recent data. If user disables Serial capture and enables it again, iDRAC starts appending from last update.

The System serial data capture starts when user enables the serial data capture flag from any of the interfaces. If serial data capture is enabled after the system has booted, you have to reboot the system, so BIOS can see the new setting (console redirection Enabled requested by iDRAC) to get the serial data. iDRAC will start the data capture continuously and stores to the shared memory with limit of 512 KB. This buffer will be circular.

() NOTE:

- For this feature to be functional, one must have Login privilege and System control privilege.
- This feature requires iDARC Datacenter license.

Dynamic configuration of virtual addresses, initiator, and storage target settings

You can dynamically view and configure the virtual address, initiator and storage target settings, and apply a persistence policy. It allows the application to apply the settings based on power state changes (that is, operating system restart, warm reset, cold reset, or AC cycle) and also based on persistence policy setting for that power state. This provides more flexibility in deployments that need rapid re-configuration of system workloads to another system.

The virtual addresses are:

- Virtual MAC Address
- Virtual iSCSI MAC Address
- Virtual FIP MAC Address
- Virtual WWN
- Virtual WWPN

NOTE: When you clear the persistence policy, all the virtual addresses are reset to the default permanent address set at the factory.

NOTE: Some cards with the virtual FIP, virtual WWN, and virtual WWPN MAC attributes, the virtual WWN and virtual WWPN MAC attributes are automatically configured when you configure virtual FIP.

Using the IO Identity feature, you can:

- View and configure the virtual addresses for network and fibre channel devices (for example, NIC, CNA, FC HBA).
- Configure the initiator (for iSCSI and FCoE) and storage target settings (for iSCSI, FCoE, and FC).
- Specify persistence or clearance of the configured values over a system AC power loss, cold, and warm system resets.

The values configured for virtual addresses, initiator and storage targets may change based on the way the main power is handled during system reset and whether the NIC, CNA, or FC HBA device has auxiliary power. The persistence of IO identity settings can be achieved based on the policy setting made using iDRAC.

Only if the I/O identity feature is enabled, the persistence policies take effect. Each time the system resets or powers on, the values are persisted or cleared based on the policy settings.

(i) NOTE: After the values are cleared, you cannot re-apply the values before running the configuration job.

Supported cards for IO Identity Optimization

The following table provides the cards that support the I/O Identity Optimization feature.

Table 43. Supported cards for I/O Identity Optimization

Manufacturer	Туре
Broadcom	• 5719 Mezz 1GB
	• 5720 PCle 1 GB

²¹⁰ Inventorying, monitoring, and configuring network devices



Table 43. Supported cards for I/O Identity Optimization (continued)

Manufacturer	Туре		
	 5720 bNDC 1 GB 5720 rNDC 1 GB 57414 PCle 25GbE 		
Intel	 i350 DP FH PCle 1GB i350 QP PCle 1GB i350 QP rNDC 1GB i350 Mezz 1GB i350 bNDC 1GB x520 PCle 10GB x520 bNDC 10GB x520 Mezz 10GB x520 Mezz 10GB x520 + i350 rNDC 10GB+1GB X710 bNDC 10GB X710 QP bNDC 10GB X710 PCle 10 GB X710 PCle 10 GB X710 rNDC 10GB+1GB X710 rNDC 10GB XL710 rNDC 10GB XL710 QSFP DP LP PCle 40GE XL710 QSFP DP FH PCle 40GE X550 DP BT PCle 2 x 10 Gb X550 DP BT LP PCle 2 x 10 Gb XXV710 Fab A/B Mezz 25 Gb (for MX platforms) 		
Mellanox	 ConnectX-3 Pro 10G Mezz 10GB ConnectX-4 LX 25GE SFP DP rNDC 25GB ConnectX-4 LX 25GE DP FH PCIe 25GB ConnectX-4 LX 25GE DP LP PCIe 25GB ConnectX-4 LX Fab A/B Mezz 25GB (for MX platforms) 		
Qlogic	 57810 PCle 10GB 57810 bNDC 10GB 57810 Mezz 10GB 57800 rNDC 10GB+1GB 57840 rNDC 10GB 57840 bNDC 10GB GME2662 Mezz FC16 QLE 2692 SP FC16 Gen 6 HBA FH PCle FC16 SP FC16 Gen 6 HBA LP PCle FC16 QLE 2690 DP FC16 Gen 6 HBA FH PCle FC16 QLE 2742 DP FC32 Gen 6 HBA FH PCle FC32 DP FC32 Gen 6 HBA LP PCle FC32 QLE2740 PCle FC32 QLE2740 PCle FC32 QME2692-DEL Fab C Mezz FC16 (for MX platforms) QL41262HMKR-DE Fab A/B Mezz 25 Gb (for MX platforms) QL41232HMKR-DE Fab A/B Mezz 25 Gb (for MX platforms) QLogic 1x32Gb QLE2770 FC HBA 		
Emulex	 LPe15002B-M8 (FH) PCIe FC8 LPe15002B-M8 (LP) PCIe FC8 LPe15000B-M8 (FH) PCIe FC8 LPe15000B-M8 (LP) PCIe FC8 		



Table 43. Supported cards for I/O Identity Optimization (continued)

Manufacturer	Туре
	LPe31000-M6-SP PCIe FC16
	LPe31002-M6-D DP PCIe FC16
	LPe32000-M2-D SP PCIe FC32
	LPe32002-M2-D DP PCIe FC32
	• LPe31002-D Fab C Mezz FC16 (for MX platforms)
	• LPe32002-D Fab C Mezz FC32 (for MX platforms)
	• LPe35002-M2 FC32 2-Port
	• LPe35000-M2 FC32 1-Port

Supported NIC firmware versions for IO Identity Optimization

In 14th generation Dell PowerEdge servers, the required NIC firmware is available by default. The following table provides the NIC firmware versions for the I/O identity optimization feature.

Virtual or Remote assigned Address and Persistence Policy behavior when iDRAC is set to Remote-Assigned Address mode or Console mode

The following table describes the Virtual Address Management (VAM) configuration and Persistence Policy behavior, and the dependencies.

Remote assigned Address Feature State in OME Modular	Mode set in iDRAC	IO Identity Feature State in iDRAC	SCP	Persistence Policy	Clear Persistence Policy — Virtual Address
Remote-Assigned Address enabled	RemoteAssignedAd dress Mode	Enabled	Virtual address management (VAM) configured	Configured VAM persists	Set to Remote assigned Address
Remote-Assigned Address enabled	RemoteAssignedAd dress Mode	Enabled	VAM not configured	Set to Remote assigned Address	No persistence — Is set to Remote assigned Address
Remote-Assigned Address enabled	RemoteAssigned Address Mode	Disabled	Configured using the path provided in Lifecycle Controller	Set to Remote assigned Address for that cycle	No persistence — Is set to Remote assigned Address
Remote-Assigned Address enabled	RemoteAssigned Address Mode	Disabled	VAM not configured	Set to Remote assigned Address	Set to Remote assigned Address
Remote-Assigned Address disabled	RemoteAssigned Address Mode	Enabled	VAM configured	Configured VAM persists	Persistence only — clear is not possible
Remote-Assigned Address disabled	RemoteAssigned Address Mode	Enabled	VAM not configured	Set to hardware MAC address	No persistence supported. Depends on card behavior
Remote-Assigned Address disabled	RemoteAssigned Address Mode	Disabled	Configured using the path provided in Lifecycle Controller	Lifecycle Controller configuration persists for that cycle	No persistence supported. Depends on card behavior
Remote-Assigned Address disabled	RemoteAssigned Address Mode	Disabled	VAM not configured	Set to hardware MAC address	Set to hardware MAC address

Table 44. Virtual/Remote-Assigned Address and Persistence Policy behavior

212 Inventorying, monitoring, and configuring network devices

Inserido ao protocolo **22.951.206-4** por: **Pedro Henrique Golin Linhares** em: 31/03/2025 12:49. A autenticidade deste documento pode ser validada no endereço: https://www.eprotocolo.pr.gov.br/spiweb/validarDocumento com o código: 692a5ef82beb173ed33dad4ee6505546.



Table 44. Virtual/Remote-Assigned Address and Persistence Policy behavior (continued)

Remote assigned Address Feature State in OME Modular	Mode set in iDRAC	IO Identity Feature State in iDRAC	SCP	Persistence Policy	Clear Persistence Policy — Virtual Address
Remote-Assigned Address enabled	Console Mode	Enabled	VAM configured	Configured VAM persists	Both persistence and clear must work
Remote-Assigned Address enabled	Console Mode	Enabled	VAM not configured	Set to hardware MAC address	Set to hardware MAC address
Remote-Assigned Address enabled	Console Mode	Disabled	Configured using the path provided in Lifecycle Controller	Lifecycle Controller configuration persists for that cycle	No persistence supported. Depends on card behavior
Remote-Assigned Address disabled	Console Mode	Enabled	VAM configured	Configured VAM persists	Both persistence and clear must work
Remote-Assigned Address disabled	Console Mode	Enabled	VAM not configured	Set to hardware MAC address	Set to hardware MAC address
Remote-Assigned Address disabled	Console Mode	Disabled	Configured using the path provided in Lifecycle Controller	Lifecycle Controller configuration persists for that cycle	No persistence supported. Depends on card behavior
Remote-Assigned Address enabled	Console Mode	Disabled	VAM not configured	Set to hardware MAC address	Set to hardware MAC address

System behavior for FlexAddress and IO Identity

Table 45. System behavior for FlexAddress and I/O Identity

Туре	FlexAddress Feature State in CMC	IO Identity Feature State in iDRAC	Availability of Remote Agent VA for the Reboot Cycle	VA Programming Source	Reboot Cycle VA Persistence Behavior
Server with FA-equivalent	Enabled	Disabled		FlexAddress from CMC	Per FlexAddress spec
Persistence	N/A, Enabled, or Disabled	Enabled	Yes - New or Persisted	Remote Agent Virtual Address	Per FlexAddress spec
			No	Virtual Address Cleared	
	Disabled	Disabled			
Server with VAM Persistence Policy Feature	Enabled	Disabled		FlexAddress from CMC	Per FlexAddress spec
	Enabled Enabled	Yes — New or Persisted	Remote Agent Virtual Address	Per Remote Agent Policy Setting	
			No	FlexAddress from CMC	Per FlexAddress spec
	Disabled	Enabled	Yes — New or Persisted	Remote Agent Virtual Address	Per Remote Agent Policy Setting
			No	Virtual Address Cleared	
	Disabled	Disabled			



Enabling or disabling IO Identity Optimization

Normally, after the system boots, the devices are configured and then after a reboot the devices are initialized. You can enable the I/O Identity Optimization feature to achieve boot optimization. If it is enabled, it sets the virtual address, initiator, and storage target attributes after the device is reset and before it is initialized, thus eliminating a second BIOS restart. The device configuration and boot operation occur in a single system start and is optimized for boot time performance.

Before enabling I/O identity optimization, make sure that:

- You have the Login, Configure, and System Control privileges.
- BIOS, iDRAC, and network cards are updated to the latest firmware.

After enabling I/O Identity Optimization feature, export the Server Configuration Profile file from iDRAC, modify the required I/O Identity attributes in the SCP file, and import the file back to iDRAC.

For the list of I/O Identity Optimization attributes that you can modify in the SCP file, see the *NIC Profile* document available at https://www.dell.com/support.

(i) NOTE: Do not modify non I/O Identity Optimization attributes.

Enabling or disabling IO Identity Optimization using web interface

To enable or disable I/O Identity Optimization:

 In the iDRAC Web interface, go to Configuration > System Settings > Hardware Settings > I/O Identity Optimization.

The I/O Identity Optimization page is displayed.

- 2. Click the I/O Identity Optimization tab, select the Enable option to enable this feature. To disable, clear this option.
- 3. Click Apply to apply the setting.

Enabling or disabling IO Identity Optimization using RACADM

To enable I/O Identity Optimization, use the command:

racadm set idrac.ioidopt.IOIDOptEnable Enabled

After enabling this feature, you must restart the system for the settings to take effect.

To disable I/O Identity Optimization, use the command:

racadm set idrac.ioidopt.IOIDOptEnable Disabled

To view the I/O Identity Optimization setting, use the command:

racadm get iDRAC.IOIDOpt

SSD Wear Threshold

iDRAC provides you the ability to configure thresholds of Remaining Rated Write Endurance for all SSD's and Available Spare of NVMe PCIe SSDs.

When SSD Remaining Rated Write Endurance and NVMe PCIe SSD Available Spare values are less than the threshold, then iDRAC logs this event in the LC log and depending on the alert type selection, iDRAC also performs Email alert, SNMP Trap, IPMI Alert, Logging in Remote Syslog, WS Eventing and OS log.

iDRAC alerts the user when the SSD Remaining Rated Write Endurance goes below the set threshold, so that the system admin can take a backup of SSD or replace it.

For only NVMe PCle SSDs, iDRAC displays **Available Spare** and provide a threshold to warn. The **Available Spare** is not available for SSDs which are connected behind PERC and HBA.



Configuring SSD Wear Threshold alert features using web interface

To configure Remaining Rated Write Endurance and Available Spare Alert Threshold using web interface:

- 1. In the iDRAC Web interface, go to Configuration > System Settings > Hardware Settings > SSD Wear Thresholds. The SSD Wear Thresholds page is displayed.
- Remaining Rated Write Endurance You can set the value between 1-99%. The default value is 10%. Alert type for this feature is SSD Wear Write Endurance and security alert is Warning as a result of threshold event..
- Available Spare Alert Threshold You can set the value between 1-99%. The default value is 10%. Alert type for this feature is SSD Wear Available Spare and security alert is Warning as a result of threshold event.

Configuring SSD Wear Threshold alert features using RACADM

To configure Remaining Rated Write Endurance, use the command:

racadm set System.Storage.RemainingRatedWriteEnduranceAlertThreshold n

, where n= 1 to 99%.

To configure Available Spare Alert Threshold, use the command:

racadm System.Storage.AvailableSpareAlertThreshold n

```
, where n= 1 to 99%.
```

Configuring persistence policy settings

Using IO identity, you can configure policies specifying the system reset and power cycle behaviors that determine the persistence or clearance of the virtual address, initiator, and storage target settings. Each individual persistence policy attribute applies to all ports and partitions of all applicable devices in the system. The device behavior changes between auxiliary powered devices and non-auxiliary powered devices.

() NOTE: The Persistence Policy feature may not work when set to default, if the VirtualAddressManagement attribute is set to FlexAddress (not for MX platforms) or RemoteAssignedAddress (for MX platforms) mode on iDRAC and if the FlexAddress or Remote-Assigned Address feature is disabled in CMC (not for MX platforms) or OME Modular (for MX platforms), ensure that you set the VirtualAddressManagement attribute to Console mode in iDRAC or enable the FlexAddress or Remote-Assigned Address feature in CMC or OME Modular.

You can configure the following persistence policies:

- Virtual Address: Auxiliary powered devices
- Virtual Address: Non-Auxiliary powered devices
- Initiator
- Storage target

Before applying the persistence policy, make sure to:

- Inventory the network hardware at least once, that is, enabled Collect System Inventory On Restart.
- Enable I/O Identity Optimization.

Events are logged to the Lifecycle Controller log when:

- I/O Identity Optimization is enabled or disabled.
- Persistence policy is changed.
- Virtual address, initiator and target values are set based on the policy. A single log entry is logged for the configured devices and the values that are set for those devices when the policy is applied.

Event actions are enabled for SNMP, email, or WS-eventing notifications. Logs are also included in the remote syslogs.



Default values for persistence policy

Table 46. Default values for persistence policy

Persistence Policy	AC Power Loss	Cold Boot	Warm Boot
Virtual Address: Auxiliary Powered Devices	Not selected	Selected	Selected
Virtual Address: Non-Auxiliary Powered Devices	Not selected	Not selected	Selected
Initiator	Selected	Selected	Selected
Storage Target	Selected	Selected	Selected

NOTE: When a persistent policy is disabled and when you perform the action to lose the virtual address, re-enabling the persistent policy does not retrieve the virtual address. You must set the virtual address again after you enable the persistent policy.

() NOTE: If there is a persistence policy in effect and the virtual addresses, initiator, or storage targets are set on a CNA-device partition, do not reset or clear the values configured for virtual addresses, initiator, and storage targets before changing the VirtualizationMode or the personality of the partition. The action is performed automatically when you disable the persistence policy. You can also use a configuration job to explicitly set the virtual address attributes to 0s and the initiator and storage targets values as defined in iSCSI initiator and storage target default values.

Configuring persistence policy settings using iDRAC web interface

To configure the persistence policy:

- In the iDRAC Web interface, go to Configuration > System Settings > Hardware Settings > I/O Identity Optimization.
- 2. Click I/O Identity Optimization tab.
- 3. In the **Persistence Policy** section, select one or more of the following for each persistence policy:
 - Warm Reset The virtual address or target settings persist when warm reset condition occurs.
 - Cold Reset The virtual address or target settings persist when cold reset conditions occur.
 - AC Power Loss The virtual address or target settings persist when AC power loss conditions occur.
- 4. Click Apply.

The persistence policies are configured.

Configuring persistence policy settings using RACADM

To set persistence policy, use the following racadm object with the **set** sub command:

- For virtual addresses, use iDRAC.IOIDOpt.VirtualAddressPersistencePolicyAuxPwrd and iDRAC.IOIDOpt.VirtualAddressPersistencePolicyNonAuxPwrd objects
- For initiator, use iDRAC.IOIDOPT.InitiatorPersistencePolicy object
- For storage targets, use **iDRAC.IOIDOpt.StorageTargetPersistencePolicy** object

For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

iSCSI initiator and storage target default values

The following tables provide the list of default values for iSCSI initiator and storage targets when the persistence policies are cleared.

Table 47. iSCSI initiator —default values

iSCSI Initiator	Default Values in IPv4 mode	Default Values in IPv6 mode
lscsilnitiatorlpAddr	0.0.0.0	

216 Inventorying, monitoring, and configuring network devices



Table 47. iSCSI initiator —default values (continued)

iSCSI Initiator	Default Values in IPv4 mode	Default Values in IPv6 mode
lscsilnitiatorlpv4Addr	0.0.0.0	0.0.0.0
lscsilnitiatorlpv6Addr		
lscsilnitiatorSubnet	0.0.0.0	0.0.0.0
IscsilnitiatorSubnetPrefix	0	0
lscsilnitiatorGateway	0.0.0.0	
lscsilnitiatorlpv4Gateway	0.0.0.0	0.0.0.0
lscsilnitiatorlpv6Gateway		
IscsilnitiatorPrimDns	0.0.0.0	::
lscsilnitiatorlpv4PrimDns	0.0.0.0	0.0.0.0
lscsilnitiatorlpv6PrimDns		::
IscsilnitiatorSecDns	0.0.0.0	
lscsilnitiatorlpv4SecDns	0.0.0.0	0.0.0.0
lscsilnitiatorlpv6SecDns	::	::
IscsilnitiatorName	Value Cleared	Value Cleared
lscsilnitiatorChapld	Value Cleared	Value Cleared
lscsilnitiatorChapPwd	Value Cleared	Value Cleared
IPVer	lpv4	lpv6

Table 48. ISCSI storage target attributes — default values

iSCSI Storage Target Attributes	Default Values in IPv4 mode	Default Values in IPv6 mode
ConnectFirstTgt	Disabled	Disabled
FirstTgtlpAddress	0.0.0.0	
FirstTgtTcpPort	3260	3260
FirstTgtBootLun	0	0
FirstTgtlscsiName	Value Cleared	Value Cleared
FirstTgtChapId	Value Cleared	Value Cleared
FirstTgtChapPwd	Value Cleared	Value Cleared
FirstTgtlpVer	Ірv4	
ConnectSecondTgt	Disabled	Disabled
SecondTgtlpAddress	0.0.0.0	

Inventorying, monitoring, and configuring network devices 217



Table 48. ISCSI storage target attributes — default values (continued)

iSCSI Storage Target Attributes	Default Values in IPv4 mode	Default Values in IPv6 mode
SecondTgtTcpPort	3260	3260
SecondTgtBootLun	0	0
SecondTgtIscsiName	Value Cleared	Value Cleared
SecondTgtChapId	Value Cleared	Value Cleared
SecondTgtChapPwd	Value Cleared	Value Cleared
SecondTgtlpVer	Ірv4	

218 Inventorying, monitoring, and configuring network devices

Inserido ao protocolo **22.951.206-4** por: **Pedro Henrique Golin Linhares** em: 31/03/2025 12:49. A autenticidade deste documento pode ser validada no endereço: https://www.eprotocolo.pr.gov.br/spiweb/validarDocumento com o código: **692a5ef82beb173ed33dad4ee6505546**.



Managing storage devices

Starting with iDRAC 3.15.15.15 release, iDRAC supports Boot Optimized Storage Solution (BOSS) controller in the 14th generation of PowerEdge servers. BOSS controllers are designed specifically for booting the operating system of the server. These controllers support limited RAID features and the configuration is staged.

Starting with iDRAC 4.30.30.30 release, iDRAC supports PERC 11, HBA 11, and BOOS 1.5 for AMD systems.

NOTE: BOSS controllers support only RAID level1.

NOTE: For BOSS Controllers, the complete VD information may not be available when both PD's are plugged-out and plugged-in back.

(i) NOTE: PERC 11 and later controllers support Hardware Root of Trust (RoT).

iDRAC has expanded its agent-free management to include direct configuration of the PERC controllers. It enables you to remotely configure the storage components attached to your system at run-time. These components include RAID and non-RAID controllers and the channels, ports, enclosures, and disks attached to them. For the PowerEdge Rx4xx/Cx4xx servers, PERC 9 and PERC 10 controllers are supported. For PowerEdge Rx5xx/Cx5xx AMD platform servers, PERC 11 is supported.

The complete storage subsystem discovery, topology, health monitoring, and configuration are accomplished in the Comprehensive Embedded Management (CEM) framework by interfacing with the internal and external PERC controllers through the MCTP protocol over I2C interface. For real-time configuration, CEM supports PERC9 controllers and above. The firmware version for PERC9 controllers must be 9.1 or later.

NOTE: The Software RAID (SWRAID) is not supported by CEM and thus is not supported in the iDRAC GUI. SWRAID can be managed using either RACADM, WSMAN or Redfish.

Using iDRAC, you can perform most of the functions that are available in OpenManage Storage Management including real-time (no reboot) configuration commands (for example, create virtual disk). You can completely configure RAID before installing the operating system.

You can configure and manage the controller functions without accessing the BIOS. These functions include configuring virtual disks and applying RAID levels and hot spares for data protection. You can initiate many other controller functions such as rebuilds and troubleshooting. You can protect your data by configuring data-redundancy or assigning hot spares.

The storage devices are:

- Controllers Most operating systems do not read and write data directly from the disks, but instead send read and write instructions to a controller. The controller is the hardware in your system that interacts directly with the disks to write and retrieve data. A controller has connectors (channels or ports) which are attached to one or more physical disks or an enclosure containing physical disks. RAID controllers can span the boundaries of the disks to create an extended amount of storage space— or a virtual disk using the capacity of more than one disk. Controllers also perform other tasks, such as initiating rebuilds, initializing disks, and more. To complete their tasks, controllers require special software known as firmware and drivers. In order to function properly, the controller must have the minimum required version of the firmware and the drivers installed. Different controllers have different characteristics in the way they read and write data and execute tasks. It is helpful to understand these features to most efficiently manage the storage.
- Physical disks or physical devices Reside within an enclosure or are attached to the controller. On a RAID controller, physical disks or devices are used to create virtual disks.
- Virtual disk It is storage created by a RAID controller from one or more physical disks. Although a virtual disk may be created from several physical disks, it is viewed by the operating system as a single disk. Depending on the RAID level used, the virtual disk may retain redundant data if there is a disk failure or have particular performance attributes. Virtual disks can only be created on a RAID controller.
- Enclosure It is attached to the system externally while the backplane and its physical disks are internal.
- Backplane It is similar to an enclosure. In a Backplane, the controller connector and physical disks are attached to the enclosure, but it does not have the management features (temperature probes, alarms, and so on) associated with external enclosures. Physical disks can be contained in an enclosure or attached to the backplane of a system.

() NOTE: In any MX chassis which contains storage sleds and compute sleds, iDRAC pertaining to any of the compute sleds in that chassis will report all storage sleds (both assigned and unassigned). If any one of the assigned or unassigned blades are in Warning or Critical heath state, the blade controller also reports the same status.



In addition to managing the physical disks contained in the enclosure, you can monitor the status of the fans, power supply, and temperature probes in an enclosure. You can hot-plug enclosures. Hot-plugging is defined as adding of a component to a system while the operating system is still running.

The physical devices connected to the controller must have the latest firmware. For the latest supported firmware, contact your service provider.

Storage events from PERC are mapped to SNMP traps and WSMan events as applicable. Any changes to the storage configurations are logged in the Lifecycle Log.

Table 49. PERC capability

PERC Capability	CEM configuration Capable Controller (PERC 9.1 or later)	CEM configuration Non-capable Controller (PERC 9.0 and lower)
Real-time	() NOTE: For PowerEdge Rx5xx/ Cx5xx servers, PERC 9, PERC 10, and PERC 11 controllers are supported.	Configuration is applied. An error message is displayed. Job creation is not successful and you cannot create real- time jobs using Web interface.
	If there is no existing pending or scheduled jobs for the controller, then configuration is applied.	
	If there are pending or scheduled jobs for that controller, then the jobs have to be cleared or you must wait for the jobs to be completed before applying the configuration at run-time. Run-time or real-time means, a reboot is not required.	
Staged	If all the set operations are staged, the configuration is staged and applied after reboot or it is applied at real-time.	Configuration is applied after reboot

Topics:

- Understanding RAID concepts
- Supported controllers
- Supported enclosures
- Summary of supported features for storage devices
- Inventorying and monitoring storage devices
- Viewing storage device topology
- Managing physical disks
- Managing virtual disks
- RAID Configuration Features
- Managing controllers
- Managing PCIe SSDs
- Managing enclosures or backplanes
- Choosing operation mode to apply settings
- Viewing and applying pending operations
- Storage devices apply operation scenarios
- Blinking or unblinking component LEDs
- Warm reboot

Understanding RAID concepts

Storage Management uses the Redundant Array of Independent Disks (RAID) technology to provide Storage Management capability. Understanding Storage Management requires an understanding of RAID concepts, as well as some familiarity with how the RAID controllers and operating system view disk space on your system.

220 Managing storage devices



What is RAID

RAID is a technology for managing the storage of data on the physical disks that reside or are attached to the system. A key aspect of RAID is the ability to span physical disks so that the combined storage capacity of multiple physical disks can be treated as a single, extended disk space. Another key aspect of RAID is the ability to maintain redundant data which can be used to restore data in the event of a disk failure. RAID uses different techniques, such as striping, mirroring, and parity, to store and reconstruct data. There are different RAID levels that use different methods for storing and reconstructing data. The RAID levels have different characteristics in terms of read/write performance, data protection, and storage capacity. Not all RAID levels maintain redundant data, which means for some RAID levels lost data cannot be restored. The RAID level you choose depends on whether your priority is performance, protection, or storage capacity.

() NOTE: The RAID Advisory Board (RAB) defines the specifications used to implement RAID. Although RAB defines the RAID levels, commercial implementation of RAID levels by different vendors may vary from the actual RAID specifications. An implementation of a particular vendor may affect the read and write performance and the degree of data redundancy.

Hardware and software RAID

RAID can be implemented with either hardware or software. A system using hardware RAID has a RAID controller that implements the RAID levels and processes data reads and writes to the physical disks. When using software RAID provided by the operating system, the operating system implements the RAID levels. For this reason, using software RAID by itself can slow the system performance. You can, however, use software RAID along with hardware RAID volumes to provide better performance and variety in the configuration of RAID volumes. For example, you can mirror a pair of hardware RAID 5 volumes across two RAID controllers to provide RAID controller redundancy.

RAID concepts

RAID uses particular techniques for writing data to disks. These techniques enable RAID to provide data redundancy or better performance. These techniques include:

- Mirroring Duplicating data from one physical disk to another physical disk. Mirroring provides data redundancy by maintaining two copies of the same data on different physical disks. If one of the disks in the mirror fails, the system can continue to operate using the unaffected disk. Both sides of the mirror contain the same data always. Either side of the mirror can act as the operational side. A mirrored RAID disk group is comparable in performance to a RAID 5 disk group in read operations but faster in write operations.
- Striping Disk striping writes data across all physical disks in a virtual disk. Each stripe consists of consecutive virtual disk data addresses that are mapped in fixed-size units to each physical disk in the virtual disk using a sequential pattern. For example, if the virtual disk includes five physical disks, the stripe writes data to physical disks one through five without repeating any of the physical disks. The amount of space consumed by a stripe is the same on each physical disk. The portion of a stripe that resides on a physical disk is a stripe element. Striping by itself does not provide data redundancy. Striping in combination with parity does provide data redundancy.
- Stripe size The total disk space consumed by a stripe not including a parity disk. For example, consider a stripe that contains 64KB of disk space and has 16KB of data residing on each disk in the stripe. In this case, the stripe size is 64KB and the stripe element size is 16KB.
- Stripe element A stripe element is the portion of a stripe that resides on a single physical disk.
- Stripe element size The amount of disk space consumed by a stripe element. For example, consider a stripe that contains 64KB of disk space and has 16KB of data residing on each disk in the stripe. In this case, the stripe element size is 16KB and the stripe size is 64KB.
- Parity Parity refers to redundant data that is maintained using an algorithm in combination with striping. When one of the striped disks fails, the data can be reconstructed from the parity information using the algorithm.
- Span A span is a RAID technique used to combine storage space from groups of physical disks into a RAID 10, 50, or 60 virtual disk.

RAID levels

Each RAID level uses some combination of mirroring, striping, and parity to provide data redundancy or improved read and write performance. For specific information on each RAID level, see Choosing raid levels.



Organizing data storage for availability and performance

RAID provides different methods or RAID levels for organizing the disk storage. Some RAID levels maintain redundant data so that you can restore data after a disk failure. Different RAID levels also entail an increase or decrease in the I/O (read and write) performance of a system.

Maintaining redundant data requires the use of additional physical disks. The possibility of a disk failure increases with an increase in the number of disks. Since the differences in I/O performance and redundancy, one RAID level may be more appropriate than another based on the applications in the operating environment and the nature of the data being stored.

When choosing a RAID level, the following performance and cost considerations apply:

- Availability or fault-tolerance Availability or fault-tolerance refers to the ability of a system to maintain operations and provide access to data even when one of its components has failed. In RAID volumes, availability or fault-tolerance is achieved by maintaining redundant data. Redundant data includes mirrors (duplicate data) and parity information (reconstructing data using an algorithm).
- Performance Read and write performance can be increased or decreased depending on the RAID level you choose. Some RAID levels may be more appropriate for particular applications.
- Cost efficiency Maintaining the redundant data or parity information associated with RAID volumes requires additional disk space. In situations where the data is temporary, easily reproduced, or non-essential, the increased cost of data redundancy may not be justified.
- Mean Time Between Failure (MTBF) Using additional disks to maintain data redundancy also increases the chance of disk failure at any given moment. Although this option cannot be avoided in situations where redundant data is a requirement, it does have implications on the workload of the system support staff within your organization.
- Volume Volume refers to a single disk non-RAID virtual disk. You can create volumes using external utilities like the O-ROM <Ctrl> <r>. Storage Management does not support the creation of volumes. However, you can view volumes and use drives from these volumes for creation of new virtual disks or Online Capacity Expansion (OCE) of existing virtual disks, provided free space is available.

Choosing RAID levels

You can use RAID to control data storage on multiple disks. Each RAID level or concatenation has different performance and data protection characteristics.

(i) NOTE: The H3xx PERC controllers do not support RAID levels 6 and 60.

The following topics provide specific information on how each RAID level store data as well as their performance and protection characteristics:

- Raid level 0 (striping)
- Raid level 1 (mirroring)
- Raid level 5 (striping with distributed parity)
- Raid level 6 (striping with additional distributed parity)
- Raid level 50 (striping over raid 5 sets)
- Raid level 60 (striping over raid 6 sets)
- Raid level 10 (striping over mirror sets)

RAID level 0 - striping

RAID 0 uses data striping, which is writing data in equal-sized segments across the physical disks. RAID 0 does not provide data redundancy.





RAID 0 characteristics:

- Groups *n* disks as one large virtual disk with a capacity of (smallest disk size) **n* disks.
- Data is stored to the disks alternately.
- No redundant data is stored. When a disk fails, the large virtual disk fails with no means of rebuilding the data.
- Better read and write performance.

RAID level 1 - mirroring

RAID 1 is the simplest form of maintaining redundant data. In RAID 1, data is mirrored or duplicated on one or more physical disks. If a physical disk fails, data can be rebuilt using the data from the other side of the mirror.



RAID 1 characteristics:

- Groups *n* + *n* disks as one virtual disk with the capacity of *n* disks. The controllers currently supported by Storage Management allow the selection of two disks when creating a RAID 1. Because these disks are mirrored, the total storage capacity is equal to one disk.
- Data is replicated on both the disks.
- When a disk fails, the virtual disk still works. The data is read from the mirror of the failed disk.
- Better read performance, but slightly slower write performance.



- Redundancy for protection of data.
- RAID 1 is more expensive in terms of disk space since twice the number of disks are used than required to store the data without redundancy.

RAID level 5 or striping with distributed parity

RAID 5 provides data redundancy by using data striping in combination with parity information. Rather than dedicating a physical disk to parity, the parity information is striped across all physical disks in the disk group.



RAID 5 characteristics:

- Groups n disks as one large virtual disk with a capacity of (n-1) disks.
- Redundant information (parity) is alternately stored on all disks.
- When a disk fails, the virtual disk still works, but it is operating in a degraded state. The data is reconstructed from the surviving disks.
- Better read performance, but slower write performance.
- Redundancy for protection of data.

RAID level 6-striping with additional distributed parity

RAID 6 provides data redundancy by using data striping in combination with parity information. Similar to RAID 5, the parity is distributed within each stripe. RAID 6, however, uses an additional physical disk to maintain parity, such that each stripe in the disk group maintains two disk blocks with parity information. The additional parity provides data protection in the event of two disk failures. In the following image, the two sets of parity information are identified as **P** and **Q**.





RAID 6 characteristics:

- Groups n disks as one large virtual disk with a capacity of (n-2) disks.
- Redundant information (parity) is alternately stored on all disks.
- The virtual disk remains functional with up to two disk failures. The data is reconstructed from the surviving disks.
- Better read performance, but slower write performance.
- Increased redundancy for protection of data.
- Two disks per span are required for parity. RAID 6 is more expensive in terms of disk space.

RAID level 50 - striping over RAID 5 sets

RAID 50 is striping over more than one span of physical disks. For example, a RAID 5 disk group that is implemented with three physical disks and then continues on with a disk group of three more physical disks would be a RAID 50.

It is possible to implement RAID 50 even when the hardware does not directly support it. In this case, you can implement more than one RAID 5 virtual disks and then convert the RAID 5 disks to dynamic disks. You can then create a dynamic volume that is spanned across all RAID 5 virtual disks.





RAID 50 characteristics:

- Groups *n***s* disks as one large virtual disk with a capacity of *s**(*n*-1) disks, where *s* is the number of spans and *n* is the number of disks within each span.
- Redundant information (parity) is alternately stored on all disks of each RAID 5 span.
- Better read performance, but slower write performance.
- Requires as much parity information as standard RAID 5.
- Data is striped across all spans. RAID 50 is more expensive in terms of disk space.

RAID level 60 - striping over RAID 6 sets

RAID 60 is striping over more than one span of physical disks that are configured as a RAID 6. For example, a RAID 6 disk group that is implemented with four physical disks and then continues on with a disk group of four more physical disks would be a RAID 60.





RAID 60 characteristics:

- Groups n^*s disks as one large virtual disk with a capacity of $s^*(n-2)$ disks, where s is the number of spans and n is the number of disks within each span.
- Redundant information (parity) is alternately stored on all disks of each RAID 6 span.
- Better read performance, but slower write performance.
- Increased redundancy provides greater data protection than a RAID 50.
- Requires proportionally as much parity information as RAID 6.
- Two disks per span are required for parity. RAID 60 is more expensive in terms of disk space.

RAID level 10 - striped with mirrors

The RAB considers RAID level 10 to be an implementation of RAID level 1. RAID 10 combines mirrored physical disks (RAID 1) with data striping (RAID 0). With RAID 10, data is striped across multiple physical disks. The striped disk group is then mirrored onto another set of physical disks. RAID 10 can be considered a *mirror of stripes*.





RAID 10 characteristics:

- Groups *n* disks as one large virtual disk with a capacity of (n/2) disks, where *n* is an even integer.
- Mirror images of the data are striped across sets of physical disks. This level provides redundancy through mirroring.
- When a disk fails, the virtual disk still works. The data is read from the surviving mirrored disk.
- Improved read performance and write performance.
- Redundancy for protection of data.

Comparing RAID level performance

The following table compares the performance characteristics associated with the more common RAID levels. This table provides general guidelines for choosing a RAID level. Evaluate your specific environment requirements before choosing a RAID level.

RAID Level	Data Redundancy	Read Performance	Write Performance	Rebuild Performance	Minimum Disks Required	Suggested Uses
RAID 0	None	Very Good	Very Good	N/A	Ν	Noncritical data.
RAID 1	Excellent	Very Good	Good	Good	2N (N = 1)	Small databases, database logs, and critical information.
RAID 5	Good	Sequential reads: good. Transactional reads: Very good	Fair, unless using writeback cache	Fair	N + 1 (N = at least two disks)	Databases and other read intensive transactional uses.
RAID 10	Excellent	Very Good	Fair	Good	2N x X	Data intensive environments (large records).

Table 50. RAID level performance comparison



Table 50. RAID level performance comparison (continued)

RAID Level	Data Redundancy	Read Performance	Write Performance	Rebuild Performance	Minimum Disks Required	Suggested Uses		
RAID 50	Good	Very Good	Fair	Fair	N + 2 (N = at least 4)	Medium sized transactional or data intensive uses.		
RAID 6	Excellent	Sequential reads: good. Transactional reads: Very good	Fair, unless using writeback cache	Poor	N + 2 (N = at least two disks)	Critical information. Databases and other read intensive transactional uses.		
RAID 60	Excellent	Very Good	Fair	Poor	X x (N + 2) (N = at least 2)	Critical information. Medium sized transactional or data intensive uses.		
N = Number of p	N = Number of physical disks							
X = Number of R	AID sets							

Supported controllers

Supported RAID controllers

The iDRAC interfaces support the following BOSS controllers:

- BOSS-S1 Adapter
- BOSS-S1 Modular (for blade servers)
- BOSS-S2 Adapter

The iDRAC interfaces support the following PERC11 controllers:

- PERC H755 Adapter
- PERC H755 Front
- PERC H755N Front
- The iDRAC interfaces support the following PERC10 controllers:
- PERC H740P Mini
- PERC H740P Adapter
- PERC H840 Adapter
- PERC H745P MX

The iDRAC interfaces support the following PERC9 controllers:

- PERC H330 Mini
- PERC H330 Adapter
- PERC H730P Mini
- PERC H730P Adapter
- PERC H730P MX

Supported non-RAID controllers

The iDRAC interface supports 12 Gbps SAS HBA external controller and HBA330 Mini or Adapter controllers.

iDRAC supports HBA330 MMZ, HBA330 MX adapters.



Supported enclosures

iDRAC supports MD1400 and MD1420 enclosures.

(i) NOTE: Redundant Array of Inexpensive Disks (RBODS) that are connected to HBA controllers are not supported.

(i) NOTE: PERC H480 with version 10.1 or greater, firmware supports up to 4 enclosures per port.

Summary of supported features for storage devices

The following tables provide the features supported by the storage devices through iDRAC.

Table 51. Supported features for storage controllers

Feature	PERC 11			PERC 10			PERC 9				
	H755 Front	H755N Front	H755 Adapter	H740P Mini	H740P Adapter	H840 Adapter	H330 Mini	H330 Adapter	H730P Mini	H730P Adapter	FD33xS
Assign or unassign physical disk as a global hot spare	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time
Convert to RAID	Not applicabl e	Not applicabl e	Not applicabl e	Not applicabl e	Not applicabl e	Not applicabl e	Not applicabl e	Not applicabl e	Not applicabl e	Not applicabl e	Not applicabl e
Convert to RAID/ Non RAID,	Real- time (convert s drive to nonRAID ePD-PT volume)	Real- time (convert s drive to nonRAID ePD-PT volume)	Real- time (convert s drive to nonRAID ePD-PT volume)	Real- time (only supporte d in eHBA controlle r mode, converts drive to nonRAID ePD-PT volume)	Real- time (only supporte d in eHBA controlle r mode, converts drive to nonRAID ePD-PT volume)	Real- time (only supporte d in eHBA controlle r mode, converts drive to nonRAID ePD-PT volume)	Real- time	Real- time	Real- time	Real- time	Real- time
Rebuild	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time
Cancel Rebuild	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time
Create virtual disks	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time
Rename virtual disks	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time
Edit virtual disks cache policies	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time

Inserido ao protocolo **22.951.206-4** por: **Pedro Henrique Golin Linhares** em: 31/03/2025 12:49. A autenticidade deste documento pode ser validada no endereço: https://www.eprotocolo.pr.gov.br/spiweb/validarDocumento com o código: 692a5ef82beb173ed33dad4ee6505546.



Table 51. Supported features for storage controllers (continued)

Feature	PERC 11			PERC 10			PERC 9				
	H755 Front	H755N Front	H755 Adapter	H740P Mini	H740P Adapter	H840 Adapter	H330 Mini	H330 Adapter	H730P Mini	H730P Adapter	FD33xS
Check virtual disk consiste ncy	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time
Cancel check consiste ncy	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time
lnitialize virtual disks	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time
Cancel initializat ion	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time
Encrypt virtual disks	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Not applicabl e	Not applicabl e	Real- time	Real- time	Real- time
Assign or unassign dedicate d hot spare	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time
Delete virtual disks	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time
Cancel Backgro und Initializat ion	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time
Online Capacity Expansio n	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time
RAID Level Migratio n	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time
Discard Preserve d Cache	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Not applicabl e	Not applicabl e	Real- time	Real- time	Real- time
Set Patrol Read Mode	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time
Manual Patrol Read Mode	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time



Table 51. Supported features for storage controllers (continued)

Feature	PERC 11			PERC 10			PERC 9				
	H755 Front	H755N Front	H755 Adapter	H740P Mini	H740P Adapter	H840 Adapter	H330 Mini	H330 Adapter	H730P Mini	H730P Adapter	FD33xS
Patrol Read Unconfig ured Areas	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time (only in web interface)				
Check Consiste ncy Mode	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time
Copybac k Mode	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time
Load Balance Mode	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time
Check Consiste ncy Rate	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time
Rebuild Rate	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time
BGI Rate	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time
Reconst ruct Rate	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time
Import foreign configur ation	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time
Auto- import foreign configur ation	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time
Clear foreign configur ation	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time
Reset controlle r configur ation	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time
Create or change security keys	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Not applicabl e	Not applicabl e	Real- time	Real- time	Real- time

Inserido ao protocolo **22.951.206-4** por: **Pedro Henrique Golin Linhares** em: 31/03/2025 12:49. A autenticidade deste documento pode ser validada no endereço: https://www.eprotocolo.pr.gov.br/spiweb/validarDocumento com o código: **692a5ef82beb173ed33dad4ee6505546**.



Table 51. Supported features for storage controllers (continued)

Feature	PERC 11			PERC 10			PERC 9				
	H755 Front	H755N Front	H755 Adapter	H740P Mini	H740P Adapter	H840 Adapter	H330 Mini	H330 Adapter	H730P Mini	H730P Adapter	FD33xS
Secure Enterpri se Key Manger	Staged	Staged	Staged	Staged	Staged	Staged	Not applicabl e	Not applicabl e	Not applicabl e	Not applicabl e	Not applicabl e
Inventor y and remotely monitor the health of PCle SSD devices	Not applicabl e										
Prepare the PCle SSD to be removed	Not applicabl e										
Securely erase the data for PCle SSD	Not applicabl e	Real- time	Not applicabl e								
Configur e Backplan e mode (split/ unified)	Real- time										
Blink or unblink compon ent LEDs	Real- time										
Switch controlle r mode	Not applicabl e	Not applicabl e	Not applicabl e	Staged							
T10PI support for Virtual Disks	Not applicabl e										

(i) NOTE: Added support for

• eHBA mode for PERC 10.2 or greater firmware which supports convert to Non-RAID disks

- convert controller to HBA mode
- RAID 10 uneven span



Table 52. Supported features of storage controllers for MX platforms

Features	PERC 11	PERC 10	PERC 9
	H755 MX	H745P MX	H730P MX
Initialize Virtual Disks	Real-time	Real-time	Real-time
Cancel Initialization	Real-time	Real-time	Real-time
Encrypt Virtual Disks	Real-time	Real-time	Real-time
Assign or unassign dedicated hot spare	Real-time	Real-time	Real-time
Delete Virtual Disks	Real-time	Real-time	Real-time
Cancel Background Initialization	Real-time	Real-time	Real-time
Online Capacity Expansion	Real-time	Real-time	Real-time
RAID Level Migration	Real-time	Real-time	Real-time
Discard Preserved Cache	Real-time	Real-time	Real-time
Set Patrol Read Mode	Real-time	Real-time	Real-time
Manual Patrol Read Mode	Real-time	Real-time	Real-time
Patrol Read Unconfigured Areas	Real-time	Real-time	Real-time (only in web interface)
Check Consistency Mode	Real-time	Real-time	Real-time
Copyback Mode	Real-time	Real-time	Real-time
Load Balance Mode	Real-time	Real-time	Real-time
Check Consistency Rate	Real-time	Real-time	Real-time
Rebuild Rate	Real-time	Real-time	Real-time
BGI Rate	Real-time	Real-time	Real-time
Reconstruct Rate	Real-time	Real-time	Real-time
Import Foreign Configuration	Real-time	Real-time	Real-time
Auto-import Foreign Configuration	Real-time	Real-time	Real-time
Clear Foreign Configuration	Real-time	Real-time	Real-time
Reset Controller Configuration	Real-time	Real-time	Real-time
Create or change security keys	Real-time	Real-time	Real-time
Inventory and remotely monitor the health of PCIe SSD devices	Real-time	Not applicable	Not applicable
Prepare the PCIe SSD to be removed	Not applicable	Not applicable	Not applicable
Securely erase the data for PCIe SSD	Real-time	Not applicable	Not applicable
Configure Backplane mode (split/ unified)	Real-time	Not applicable	Not applicable
Blink or unblink component LEDs	Real-time	Real-time	Real-time

Inserido ao protocolo **22.951.206-4** por: **Pedro Henrique Golin Linhares** em: 31/03/2025 12:49. A autenticidade deste documento pode ser validada no endereço: https://www.eprotocolo.pr.gov.br/spiweb/validarDocumento com o código: **692a5ef82beb173ed33dad4ee6505546**.



Table 52. Supported features of storage controllers for MX platforms (continued)

Features	PERC 11	PERC 10	PERC 9	
	H755 MX	H745P MX	H730P MX	
Switch Controller Mode	Not applicable	Not applicable	Staged	
T10PI Support for Virtual Disks	Not applicable	Not applicable	Not applicable	

(i) NOTE: H745P MX supports eHBA mode with PERC 10.2 and higher.

Table 53. Supported features for storage devices

Feature	PCIe SSD	BOSS S1	BOSS S2
Create Virtual Disks	Not applicable	Staged	Staged
Reset Controller Configuration	Not applicable	Staged	Staged
Fast Initialization	Not applicable	Staged	Staged
Delete Virtual Disks	Not applicable	Staged	Staged
Full Initialization	Not applicable	Not applicable	Not applicable
Inventory and remotely monitor the health of PCIe SSD devices	Real-time	Not applicable	Not applicable
Prepare the PCIe SSD to be removed	Real-time	Not applicable	Not applicable
Securely erase the data for PCIe SSD	Staged	Not applicable	Not applicable
Blink or unblink component LEDs	Real-time	Not applicable	Real-time
Hot plugging of drives	Real-time	Not applicable	Real-time

Inventorying and monitoring storage devices

You can remotely monitor the health and view the inventory of the following Comprehensive Embedded Management (CEM) enabled storage devices in the managed system using iDRAC web interface:

- RAID controllers, non-RAID controllers, BOSS controllers and PCIe extenders
- Enclosures that include Enclosure Management Modules (EMMs), power supply, fan probe, and temperature probe
- Physical disks
- Virtual disks
- Batteries

The recent storage events and topology of storage devices are also displayed.

Alerts and SNMP traps are generated for storage events. The events are logged in the Lifecycle Log.

() NOTE:

- If you enumerate the enclosure view's WSMan command on a system while one PSU-cable is removed, the primary status of the enclosure view is reported as **Healthy** instead of **Warning**.
- For an accurate inventory of BOSS controllers, ensure that Collect System Inventory On Reboot Operation (CSIOR) is completed. CSIOR is enabled by default.
- The storage health rollup follows the same convention of Dell EMC OpenManage product. For more information see the *OpenManage Server Administrator User's Guide* available at https://www.dell.com/openmanagemanuals.



- Physical disks in system with multiple backplanes may be listed under a different backplane. Use the blink function to identify the disks.
- FQDD of certain Backplanes may not be same in Software Inventory and Hardware Inventory.
- Lifecycle log for PERC controller is not available when the past PERC controller events are being processed and this does not affect the functionality. Past event processing can vary depending on the configuration

Monitoring storage devices using web interface

To view the storage device information using web interface:

- Go to **Storage** > **Overview** > **Summary** to view the summary of the storage components and the recently logged events. This page is automatically refreshed every 30 seconds.
- Go to **Storage** > **Overview** > **Controllers** to view the RAID controller information. The **Controllers** page is displayed.
- Go to **Storage** > **Overview** > **Physical Disks** to view physical disk information. The **Physical Disks** page is displayed.
- Go to Storage > Overview > Virtual Disks to view virtual disk information. The Virtual Disks page is displayed.
- Go to Storage > Overview > Enclosures to view the enclosure information. The Enclosures page is displayed.

You can also use filters to view specific device information.

() NOTE:

- The storage hardware list is not displayed in case the system does not have storage devices with CEM support.
- Behavior of non-Dell certified or 3rd party NVMe devices may not be consistent in iDRAC.
- If the NVMe SSDs in the backplane slots support NVMe-MI commands and the I2C connection to backplane slots are fine, the iDRAC discovers these NVMe SSDs and reports them in the interfaces irrespective of the PCI connections to the respective backplane slots.

(i) NOTE:

Туре	Web GUI Support	Other Interfaces support
SATA	Not available	Inventory and RAID configuration
NVMe	Physical disk inventory only	Inventory and RAID configuration

For more information about the displayed properties and to use the filter options, see the iDRAC Online Help.

Monitoring storage devices using RACADM

To view the storage device information, use the storage command.

For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Monitoring backplane using iDRAC settings utility

In the iDRAC Settings utility, go to **System Summary**. The **iDRAC Settings.System Summary** page is displayed. The **Backplane Inventory** section displays the backplane information. For information about the fields, see the *iDRAC Settings Utility Online Help*.

Viewing storage device topology

You can view the hierarchical physical containment view of the key storage components, that is, a list of controllers, enclosures connected to the controller and a link to the physical disk contained in each enclosure. The physical disks attached directly to the controller are also displayed.

To view the storage device topology, go to **Storage** > **Overview**. The **Overview** page displays the hierarchical representation of the storage components in the system. The available options are:

- Controllers
- Physical Disks

236 Managing storage devices



- Virtual Disks
- Enclosures

Click the links to view the respective component details.

Managing physical disks

You can perform the following for physical disks:

- View physical disk properties.
- Assign or unassign physical disk as a global hot-spare.
- Convert to RAID capable disk.
- Convert to non-RAID disk.
- Blink or unblink the LED.
- Rebuild physical disk
- Cancel rebuild physical disk
- Cryptographic erase

Assigning or unassigning physical disk as global hot spare

A global hot spare is an unused backup disk that is part of the disk group. Hot spares remain in standby mode. When a physical disk that is used in a virtual disk fails, the assigned hot spare is activated to replace the failed physical disk without interrupting the system or requiring your intervention. When a hot spare is activated, it rebuilds the data for all redundant virtual disks that were using the failed physical disk.

(i) NOTE: From iDRAC v3.00.00.00 or later, you can add global hot spares when virtual disks are not created.

You can change the hot spare assignment by unassigning a disk and choosing another disk as needed. You can also assign more than one physical disk as a global hot spare.

Global hot spares must be assigned and unassigned manually. They are not assigned to specific virtual disks. If you want to assign a hot spare to a virtual disk (it replaces any physical disk that fails in the virtual disk), then see Assigning or unassigning dedicated hot spares.

When deleting virtual disks, all assigned global hot spares may be automatically unassigned when the last virtual disk associated with the controller is deleted.

If you reset the configuration, the virtual disks are deleted and all the hot spares are unassigned.

You must be familiar with the size requirements and other considerations associated with hot spares.

Before assigning a physical disk as a global hot spare:

- Make sure that Lifecycle Controller is enabled.
- If there are no disk drives available in ready state, insert additional disk drives and make sure that the drives are in ready state.
- If physical disks are in non-RAID mode convert them to RAID mode using iDRAC interfaces such as iDRAC web interface, RACADM, Redfish or WSMan, or <CTRL+R>.
 - **NOTE:** During POST, press F2 to enter System Setup or Device Setup. CTRL+R option is no longer supported for PERC 10. CTRL+R only works with PERC 9 while boot mode is set to BIOS.

If you have assigned a physical disk as a global hot spare in Add to Pending Operation mode, the pending operation is created but a job is not created. Then, if you try to unassign the same disk as global hot spare, the assign global hot spare pending operation is cleared.

If you have unassigned a physical disk as a global hot spare in Add to Pending Operation mode, the pending operation is created but a job is not created. Then, if you try to assign the same disk as a global hot spare, the unassign global hot spare pending operation is cleared.

If the last VD is deleted, the global hot spares also returns to ready state.

If a PD is already a global hot spares, user can still assign it again as a global hot spares.



Assigning or unassigning global hot spare using web interface

To assign or unassign a global hot spare for a physical disk drive:

- 1. In the iDRAC web interface, go to **Configuration** > **Storage Configuration**.
- The Storage Configuration page is displayed.
- 2. From the **Controller** drop-down menu, select the controller to view the associated physical disks.
- **3.** Click **Physical Disk Configuration**. All the physical disks associated to the controller are displayed.
- 4. To assign as a global hotspare, from the drop-down menus in the **Action** column, select **Assign Global Hotspare** for one or more physical disks.
- 5. To unassign a hot spare, from the drop-down menus in the **Action** column, select **Unassign Hotspare** for one or more physical disks.
- 6. Click Apply Now.

Depending on your requirement, you can also choose to apply **At Next Reboot** or **At Scheduled Time**. Based on the selected operation mode, the settings are applied.

Assigning or unassigning global hot spare using RACADM

Use the storage command and specify the type as global hot spare.

For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Converting a physical disk to RAID or non-RAID mode

Converting a physical disk to RAID mode enables the disk for all RAID operations. When a disk is in a non-RAID mode, the disk is exposed to the operating system unlike unconfigured good disks and is used in a direct pass-through mode.

PERC 10 is not supported to convert drives to non-RAID. But it is supported in PERC 10.2 and higher versions.

You can convert the physical disk drives to RAID or non-RAID mode by:

- Using iDRAC interfaces such as iDRAC web interface, RACADM, Redfish or WSMan.
- Pressing <Ctrl+R> while restarting the server and selecting the required controller.
- () NOTE: If the physical drives connected to a PERC controller are in non-RAID mode, the size of the disk displayed in the iDRAC interfaces, such as iDRAC GUI, RACADM, Redfish and WSMan, may be slightly less than the actual size of the disk. However, you can use the full capacity of the disk to deploy operating systems.

(i) NOTE:

- Hot plugged disks in PERC H330 are always in non-RAID mode. In other RAID controllers, they are always in RAID mode.
- Hot plugged disks in PERC 11 are either ready or EPD-PT depending on the current auto configure behavior setting.

Converting physical disks to RAID capable or non-RAID mode using the iDRAC web interface

To convert the physical disks to RAID mode or non-RAID mode, perform the following steps:

- 1. In the iDRAC web interface, click Storage > Overview > Physical Disks.
- Click Filter options. Two options are displayed Clear All Filters and Advanced Filter. Click Advanced Filter option. An elaborated list is displayed that allows you to configure different parameters.
- **3.** From the **Group By** drop-down menu, select an enclosure or virtual disks. The parameters associated with the enclosure or the VD are displayed.
- **4.** Click **Apply**, once you select all the desired parameters. For more information about the fields, see the *iDRAC Online Help*. The settings are applied based on the option selected in the operation mode.

Converting physical disks to RAID capable or non-RAID mode using RACADM

Depending on whether you want to convert to RAID or Non-RAID mode, use the following RACADM commands



- To convert to RAID mode, use the racadm storage converttoraid command.
- To convert to Non-RAID mode, use the racadm storage convertiononraid command.

NOTE: On the S140 controller, you can only use the RACADM interface to convert the drives from non-RAID to RAID mode. The supported Software RAID modes are Windows or Linux Mode.

For more information about the commands, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Erasing physical disks

The System Erase feature allows you to erase the contents of the physical drives. This feature is accessible using RACADM or the LC GUI. Physical drives on the server are grouped into two categories.

- Secure erase drives— Includes drives that provide cryptographic erase such as ISE and SED SAS and SATA drives, and PCIe SSDs.
- Overwrite erase drives— Includes all drives that do not support cryptographic erase.

(i) NOTE: Before erasing vFlash, you must first detach all partitions using iDRAC interfaces before executing the operation.

NOTE: System erase only applies to drives within the server. iDRAC is not able to erase drives in an external enclosure such as a JBOD.

The RACADM SystemErase sub-command includes options for the following categories:

- The SecureErasePD option cryptographically erases all the secure erase drives.
- The OverwritePD option overwrites data on all drives.
- **NOTE:** Cryptographic Erase of BOSS physical disk can be done by SystemErase method and it is supported from LC-UI, Wsman, and Racadm

Before performing SystemErase, use the following command to check the erase capability of all physical disks for a server:

racadm storage get pdisks -o -p SystemEraseCapability

() NOTE: If SEKM is enabled on the server, then disable SEKM using the racadm sekm disable command before using this command. This can avoid any storage devices being locked out which are secured by iDRAC, if SEKM settings are erased from iDRAC by executing this command.

To erase ISE and SED drives, use this command:

racadm systemerase -secureerasepd

To erase overwrite erase drives, use the following command:

racadm systemerase -overwritepd

(i) NOTE: RACADM SystemErase removes all the virtual disks from the physical disks that are erased by the above commands.

(i) NOTE: RACADM SystemErase causes the server to restart in order to perform the erase operations.

NOTE: Individual PCIe SSD or SED devices can be erased using the iDRAC GUI or RACADM. For more information, see the *Erasing PCIe SSD device data* and the *Erasing SED device data* section.

For information on the System Erase function within the Lifecycle Controller GUI, see the *Lifecycle Controller User's Guide* available at https://www.dell.com/idracmanuals .

Erasing SED/ISE device data

NOTE: This operation is not supported when supported device is a part of a Virtual Disk. The target supported device must be removed from the virtual disk prior to performing device erase.

Cryptographic Erase permanently erases all data present on the disk. Performing a Cryptographic Erase on an SED/ISE overwrites all blocks and results in permanent loss of all data on the supported devices. During Cryptographic Erase, the host is

unable to access the supported device. SED/ISE device erase can be performed either in real time or be applied after a system reboot.

Fls. <u>626</u> Mov. 32

If the system reboot or experiences a power loss during cryptographic erase, the operation is canceled. You must reboot the system and restart the process.

Before erasing SED/ISE device data, ensure that:

- Lifecycle Controller is enabled.
- You have Server Control and Login privileges.
- Selected supported drive is not part of a virtual disk.

() NOTE:

- Erasing SED/ISE can be performed either as a real time or as a staged operation.
- After the drive is erased, it may still be displayed as active within the OS due to data caching. If this occurs, reboot the OS and the erased drive will no longer be displayed or report any data.
- Cryptographic erase operation is not supported for hot-plugged NVMe disks. Reboot the server before starting the operation. If the operation continues to fail, ensure that CSIOR is enabled and that the NVMe disks qualified by Dell Technologies.

Erasing SED/ISE device data using web interface

To erase the data on the supported device:

- 1. In the iDRAC Web interface, go to **Storage** > **Overview** > **Physical Disks**.
- The **Physical Disk** page is displayed.
- 2. From the **Controller** drop-down menu, select the controller to view the associated devices.
- 3. From the drop-down menus, select Cryptographic Erase for one or more SED/ISEs.

If you have selected **Cryptographic Erase** and you want to view the other options in the drop-down menu, then select **Action** and then click the drop-down menu to view the other options.

- 4. From the Apply Operation Mode drop-down menu, select one of the following options:
 - Apply Now Select this option to apply the actions immediately with no system reboot required.
 - At Next Reboot Select this option to apply the actions during the next system reboot.
 - At Scheduled Time Select this option to apply the actions at a scheduled day and time:
 - **Start Time** and **End Time** Click the calendar icons and select the days. From the drop-down menus, select the time. The action is applied between the start time and end time.
 - From the drop-down menu, select the type of reboot:
 - No Reboot (Manually Reboot System)
 - Graceful Shutdown
 - Force Shutdown
 - Power Cycle System (cold boot)
- 5. Click Apply.

If the job is not created, a message indicating that the job creation was not successful is displayed. Also, the message ID and the recommended response action is displayed.

If the job is created successfully, a message indicating that the job ID is created for the selected controller is displayed. Click **Job Queue** to view the progress of the job in the Job Queue page.

If pending operation is not created, an error message is displayed. If pending operation is successful and job creation is not successful, then an error message is displayed.

Erasing SED device data using RACADM

To securely erase an SED device:

```
racadm storage cryptographicerase:<SED FQDD>
```

Fls. <u>627</u> Mov. <u>32</u>

To create the target job after executing the cryptographicerase command:

racadm jobqueue create <SED FQDD> -s TIME NOW -realtime

To create the target staged job after executing the cryptographicerase command:

racadm jobqueue create <SED FQDD> -s TIME NOW -e <start time>

To query the job ID returned:

racadm jobqueue view -i <job ID>

For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Rebuild Physical Disk

Rebuild Physical Disk is the ability to reconstruct the contents of a failed disk. This is true only when auto rebuild option is set to false. If there is a redundant virtual disk, the rebuild operation can reconstruct the contents of a failed physical disk. A rebuild can take place during normal operation, but it degrades performance.

Cancel Rebuild can be used to cancel a rebuild that is in progress. If you cancel a rebuild, the virtual disk remains in a degraded state. The failure of an additional physical disk can cause the virtual disk to fail and may result in data loss. It is recommended to perform a rebuild on the failed physical disk at the earliest.

In case, you cancel the rebuild of a physical disk that is assigned as a hot spare, reinitiate the rebuild on the same physical disk in order to restore the data. Canceling the rebuild of a physical disk and then assigning another physical disk as a hot spare does not cause the newly assigned hot spare to rebuild the data.

Managing virtual disks

You can perform the following operations for the virtual disks:

- Create
- Delete
- Edit policies
- Initialize
- Check consistency
- Cancel check consistency
- Encrypt virtual disks
- Assign or unassign dedicated hot spares
- Blink and unblink virtual disk
- Cancel background initialization
- Online capacity expansion
- RAID level migration

NOTE: You can manage and monitor 240 virtual disks using iDRAC interfaces. To create VDs, use either Device Setup (F2), PERCCLI command line tool, or Dell OpenManage Server Administrator (OMSA).

(i) NOTE: PERC 10 count is less since it does not support daisy chain arrangements.

Creating virtual disks

To implement RAID functions, you must create a virtual disk. A virtual disk refers to storage created by a RAID controller from one or more physical disks. Although a virtual disk may be created from several physical disks, it is seen by the operating system as a single disk.

Before creating a virtual disk, you should be familiar with the information in Considerations Before Creating Virtual Disks.

You can create a Virtual Disk using the Physical Disks attached to the PERC controller. To create a Virtual Disk, you must have the Server Control user privilege. You can create a maximum of 64 virtual drives and a maximum of 16 virtual drives in the same drive group.



You cannot create a virtual disk if:

- Physical disk drives are not available for virtual disk creation. Install additional physical disk drives.
- Maximum number of virtual disks that can be created on the controller has been reached. You must delete at least one virtual disk and then create a new virtual disk.
- Maximum number of virtual disks supported by a drive group has been reached. You must delete one virtual disk from the selected group and then create a new virtual disk.
- A job is currently running or scheduled on the selected controller. You must wait for this job to complete or you can delete the job before attempting a new operation. You can view and manage the status of the scheduled job in the Job Queue page.
- Physical disk is in non-RAID mode. You must convert to RAID mode using iDRAC interfaces such as iDRAC web interface, RACADM, Redfish, WSMan, or <CTRL+R>.

NOTE: If you create a virtual disk in Add to Pending Operation mode and a job is not created, and then if you delete the Virtual disk, then the create pending operation for the virtual disk is cleared.

(i) NOTE: RAID 6 and 60 are not supported in PERC H330.

NOTE: BOSS controller allows you to create virtual disk only of size equal to the full size of the M.2 physical storage media. Ensure that you set the virtual disk size to zero when using the Server Configuration Profile to create a BOSS virtual disk. For other interfaces such as RACADM, WSMan, and Redfish, the virtual disk size should not be specified.

Considerations before creating virtual disks

Before creating virtual disks, consider the following:

- Virtual disk names not stored on controller—The names of the virtual disks that you create are not stored on the controller. This means that if you reboot using a different operating system, the new operating system may rename the virtual disk using its own naming conventions.
- Disk grouping is a logical grouping of disks attached to a RAID controller on which one or more virtual disks are created, such that all virtual disks in the disk group use all of the physical disks in the disk group. The current implementation supports the blocking of mixed disk groups during the creation of logical devices.
- Physical disks are bound to disk groups. Therefore, there is no RAID level mixing on one disk group.
- There are limitations on the number of physical disks that can be included in the virtual disk. These limitations depend on the controller. When creating a virtual disk, controllers support a certain number of stripes and spans (methods for combining the storage on physical disks). Because the number of total stripes and spans is limited, the number of physical disks that can be used is also limited. The limitations on stripes and spans affect the RAID levels as follows:
 - Maximum number of spans affects RAID 10, RAID 50, and RAID 60.
 - Maximum number of stripes affects RAID 0, RAID 5, RAID 50, RAID 6, and RAID 60.
 - \circ $\,$ Number of physical disks in a mirror is always 2. This affects RAID 1 and RAID 10.
 - (i) NOTE:
 - RAID 1 is only supported for BOSS controllers.
 - SWRAID controller only supports RAID 0, 1, 5 and 10.
- Cannot create virtual disks on PCIe SSDs. But PERC 11 and later controllers support creating virtual disks using PCIe SSDs.

Creating virtual disks using web interface

To create virtual disk:

- 1. In the iDRAC Web interface, go to **Storage** > **Overview** > **Virtual DisksAdvanced Filter**.
- 2. In the Virtual Disk section, do the following:
 - a. From the Controller drop-down menu, select the controller for which you want to create the virtual disk.
 - b. From the Layout drop-down menu, select the RAID level for the Virtual Disk.

Only those RAID levels supported by the controller appear in the drop-down menu and it is based on the RAID levels are available based on the total number of physical disks available.

- c. Select the Media Type, Stripe Size, Read Policy, Write Policy, Disk Cache Policy, .
- Only those values supported by the controller appear in the drop-down menus for these properties.
- d. In the Capacity field, enter the size of the virtual disk. The maximum size is displayed and then updated as disks are selected.
- e. The **Span Count** field is displayed based on the selected physical disks (step 3). You cannot set this value. It is automatically calculated after selecting disks for multi-raid level. **Span Count** field is applicable to RAID 10, RAID 50, and



RAID 60. If you have selected RAID 10 and if the controller supports uneven RAID 10, then the span count value is not displayed. The controller automatically sets the appropriate value. For RAID 50 and RAID 60, this field is not displayed when minimum number of disks are used to create RAID. It can be changed if more disks are used.

- **3.** In the **Select Physical Disks** section, select the number of physical disks. For more information about the fields, see the *iDRAC Online Help*
- 4. From the Apply Operation Mode drop-down menu, select when you want to apply the settings.
- 5. Click Create Virtual Disk.

Based on the selected **Apply Operation Mode**, the settings are applied.

NOTE: You can use alphanumeric characters, spaces, dashes, and underscores in the disk name.

Any other special characters that you enter are removed and replaced by space while creating the virtual disk.

Creating virtual disks using RACADM

Use the racadm storage createvd command.

For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

(i) NOTE: Disk slicing or configuring partial VDs is not supported using RACADM on the drives managed by S140 controller.

Editing virtual disk cache policies

You can change the read, write, or disk cache policy of a virtual disk.

NOTE: Some of the controllers do not support all read or write policies. Therefore, when a policy is applied, an error message is displayed.

The read policies indicate whether the controller must read sequential sectors of the virtual disk searching for data:

- Adaptive Read Ahead The controller initiates read ahead only if the two most recent reads requests accessed sequential sectors of the disk. If subsequent read requests access random sectors of the disk, the controller reverts to no read ahead policy. The controller continues to evaluate whether read requests are accessing sequential sectors of the disk, and initiates read ahead if necessary.
- **Read Ahead** The controller reads sequential sectors of the virtual disk when seeking data. Read ahead policy may improve system performance if the data is written to the sequential sectors of the virtual disk.
- No Read Ahead Selecting no read ahead policy indicates that the controller should not use read ahead policy.

The write policies specify if the controller sends a write-request completion signal when the data is in the cache or after it has been written to the disk.

- Write Through The controller sends a write-request completion signal only after the data is written to the disk. Write-through caching provides better data security than write-back caching, since the system assumes that the data is available only after it has been safely written to the disk.
- Write Back The controller sends a write-request completion signal as soon as the data is in the controller cache but has not yet been written to disk. Write back caching may provide improved performance since subsequent read requests can retrieve data quickly from the cache then from the disk. However, data loss may occur in the event of a system failure which prevents that data from being written on a disk. Other applications may also experience problems when actions assume that the data is available on the disk.
- Force Write Back The write cache is enabled regardless of whether the controller has a battery. If the controller does not have a battery and force write-back caching is used, data loss may occur in the event of a power failure.

The Disk Cache policy applies to readings on a specific virtual disk. These settings do not affect the read-ahead policy.

() NOTE:

- Controller non-volatile cache and battery backup of controller cache affects the read-policy or the write policy that a controller can support. All PERCs do not have battery and cache.
- Read ahead and write back requires cache. Therefore, if the controller does not have cache, it does not allow you to set the policy value.

Similarly, if the PERC has cache but not battery and the policy is set that requires accessing cache, then data loss may occur if base of power off. So few PERCs may not allow that policy.

Therefore, depending upon the PERC, the policy value is set.



Deleting virtual disks

Deleting a virtual disk destroys all information including file systems and volumes residing on the virtual disk and removes the virtual disk from the controller's configuration. When deleting virtual disks, all assigned global hot spares may be automatically unassigned when the last virtual disk associated with the controller is deleted. When deleting the last virtual disk of a disk group, all assigned dedicated hot spares automatically become global hot spares.

If you delete all the VDs for a global hotspare, then the global hotspare gets automatically deleted.

You must have the Login and Server Control privilege to perform delete virtual disks.

When this operation is allowed, you can delete a boot virtual drive. It is done from sideband and the independent of the operating system. Hence, a warning message appears before you delete the virtual drive.

If you delete a virtual disk and immediately create a new virtual disk with all the same characteristics as the one that was deleted, the controller recognizes the data as if the first virtual disk were never deleted. In this situation, if you do not want the old data after recreating a new virtual disk, re-initialize the virtual disk.

Checking virtual disk consistency

This operation verifies the accuracy of the redundant (parity) information. This task only applies to redundant virtual disks. When necessary, the check consistency task rebuilds the redundant data. If the virtual drive has a degraded status, running a check consistency may be able to return the virtual drive to ready status. You can perform a consistency check using the web interface or RACADM.

You can also cancel the check consistency operation. The cancel check consistency is a real-time operation.

You must have Login and Server Control privilege to check consistency of virtual disks.

(i) NOTE: Consistency check is not supported when the drives are set up in RAIDO mode.

NOTE: If you perform Cancel Consistency operation when there is no consistency check operation is in progress, then the pending operation in GUI is shown as Cancel BGI instead of Cancel Consistency check.

Initializing virtual disks

Initializing virtual disks erases the all the data on the disk but does not change the virtual disk configuration. You must initialize a virtual disk that is configured before it is used.

(i) NOTE: Do not initialize virtual disks when attempting to recreate an existing configuration.

You can perform a fast initialization, a full Initialization, or cancel the initialization operation.

NOTE: The cancel initialization is a real-time operation. You can cancel the initialization using only the iDRAC Web interface and not RACADM.

Fast initialization

The fast initialize operation initializes all physical disks included in the virtual disk. It updates the metadata on the physical disks so that all disk space is available for future write operations. The initialize task can be completed quickly because the existing information on the physical disks is not erased, although future write operations overwrite any information that remains on the physical disks.

Fast initialization only deletes the boot sector and stripe information. Perform a fast initialize only if you are constrained for time or the hard drives are new or unused. Fast Initialization takes less time to complete (usually 30-60 seconds).

CAUTION: Performing a fast initialize causes existing data to be inaccessible.

The fast initialize task does not write zeroes to the disk blocks on the physical disks. It is because the Fast Initialize task does not perform a write operation, it causes less degradation to the disk.

244 Managing storage devices


A fast initialization on a virtual disk overwrites the first and last 8 MB of the virtual disk, clearing any boot records or partition information. The operation takes only 2-3 seconds to complete and is recommended when you are recreating virtual disks.

A background initialization starts five minutes after the Fast Initialization is completed.

Full or slow initialization

The full initialization (also called slow initialize) operation initializes all physical disks included in the virtual disk. It updates the metadata on the physical disks and erases all existing data and file systems. You can perform a full initialization after creating the virtual disk. In comparison with the fast initialize operation, you may want to use the full initialize if you have trouble with a physical disk or suspect that it has bad disk blocks. The full initialize operation remaps bad blocks and writes zeroes to all disk blocks.

If full initialization of a virtual disk is performed, background initialization is not required. During full initialization, the host is not able to access the virtual disk. If the system reboots during a full initialization, the operation terminates and a background initialization process starts on the virtual disk.

It is always recommended to do a full initialization on drives that previously contained data. Full initialization can take up to 1-2 minutes per GB. The speed of initialization depends on the controller model, speed of hard drives, and the firmware version.

The full initialize task initializes one physical disk at a time.

(i) NOTE: Full initialize is supported only in real-time. Only few controllers support full initialization.

Encrypting virtual disks

When encryption is disabled on a controller (that is, the security key is deleted), manually enable encryption for virtual disks created using SED drives. If the virtual disk is created after encryption is enabled on a controller, the virtual disk is automatically encrypted. It is automatically configured as an encrypted virtual disk unless the enabled encryption option is disabled during the virtual disk creation.

You must have Login and Server Control privilege to manage the encryption keys.

NOTE: Though encryption is enabled in the controllers, user needs to manually enable encryption on the VD if VD is created from iDRAC. Only if the VD is created from OMSA, it would be automatically encrypted.

Assigning or unassigning dedicated hot spares

A dedicated hot spare is an unused backup disk that is assigned to a virtual disk. When a physical disk in the virtual disk fails, the hot spare is activated to replace the failed physical disk without interrupting the system or requiring your intervention.

You must have Login and Server Control privilege to run this operation.

You can assign only 4K drives as hot spare to 4K virtual disks.

If you have assigned a physical disk as a dedicated hot spare in Add to Pending Operation mode, the pending operation is created but a job is not created. Then, if you try to unassign the dedicated hot spare, the assign dedicated hot spare pending operation is cleared.

If you have unassigned a physical disk as a dedicated hot spare in Add to Pending Operation mode, the pending operation is created but a job is not created. Then, if you try to assign the dedicated hot spare, the unassign dedicated hot spare pending operation is cleared.

NOTE: While the log export operation is in progress, you cannot view information about dedicated hot spares on the

Manage Virtual Disks page. After the log export operation is complete, reload or refresh the **Manage Virtual Disks** page to view the information.

Rename VD

To change the name of a Virtual Disk, the user must have System Control privilege. The virtual disk name can contain only alphanumeric characters, spaces, dashes and underscores. The maximum length of the name depends on the individual controller. In most cases, the maximum length is 15 characters. The name cannot start with a space, end with a space, or be left blank. Every time a virtual disk is renamed, an LC Log gets created.



Edit Disk capacity

Online Capacity Expansion (OCE) allows you to increase the storage capacity of selected RAID levels while the system remains online. The controller redistributes the data on the array(called Reconfiguration), placing new space available at the end of each RAID array.

Online Capacity Expansion (OCE) can be achieved in two ways:

- If free space is available on the smallest physical drive on the virtual disks group after starting LBA of Virtual disks, the virtual disk's capacity can be expanded within that free space. This option allows you to enter the new increased virtual disk size. If disk group in a virtual disk has space available only before starting LBA, then Edit Disk Capacity in same disk group is not permitted even though there is Available Space on a physical drive.
- A virtual disk's capacity can also be expanded by adding additional compatible physical disks to the existing virtual disk group. This option does not allow you to enter the new increased virtual disk size. New increased virtual disk size is calculated and displayed to the user based on the used disk space of existing physical disk group on a particular virtual disk, existing raid level of the virtual disk and the number of new drives added to the virtual disk.

Capacity Expansion allows user to specify the final VD size. Internally final VD size is conveyed to PERC in percentage (this percentage is the space user would like to use from empty space left in the array for the local disk to expand). Because of this percentage logic final VD size after reconfiguration completes may be different from what user provided for scenario where user is not giving maximum VD size possible as the final VD size (percentage turns out to be less than 100%). User does not see difference in this entered VD size and final VD size after reconfiguration, if maximum possible VD size is entered by user.

Raid Level Migration

RAID Level Migration (RLM) refers to changing a virtual disk's RAID level. iDRAC9 provides an option to increase the VD size using RLM. In a way, RLM allows migrating the RAID level of a virtual disk which in turn may increase the size of virtual disk.

RAID level migration is the process of converting a VD with one RAID Level to another. When you migrate a VD to a different Raid Level, the user data on it is redistributed to the format of the new configuration.

This configuration is supported by both staged and realtime.

The below table describes possible reconfigurable VD layouts while reconfiguring (RLM) a VD with addition of disks and without addition of disks.

Table 54. Possible VD Layout

Source VD Layout	Possible target VD Layout with Disk Add	Possible target VD Layout Without disk addition
R0 (single disk)	R1	NA
RO	R5/R6	NA
R1	R0/R5/R6	RO
R5	R0/R6	RO
R6	R0/R5	R0/R5

Permitted operations when OCE or RLM is going on

The following operations are allowed when OCE/RLM is going on:

Table 55. Permitted operations

From Controller End behind which a VD is going through OCE/RLM	From VD End (which is going through OCE/RLM)	From any other Ready State Physical Disk on the same controller	From any other VD (which is not going through OCE/ RLM) End on the same controller
Reset Configuration	Delete	Blink	Delete
Export Log	Blink	Unblink	Blink



Table 55. Permitted operations (continued)

From Controller End behind which a VD is going through OCE/RLM	From VD End (which is going through OCE/RLM)	From any other Ready State Physical Disk on the same controller	From any other VD (which is not going through OCE/ RLM) End on the same controller
Set Patrol Read Mode	Unblink	Assign Global Hot Spare	Unblink
Start Patrol Read		Convert to non-RAID Disks	Rename
Change Controller Properties			Change Policy
Manage Physical Disk Power			Slow Initialize
Convert to RAID Capable Disks			Fast Initialize
Convert to Non-RAID Disks			Replace Member Disk
Change Controller Mode			

OCE and RLM Restrictions or Limitations

Following are the common limitations for OCE and RLM:

- OCE/RLM is restricted to the scenario where the disk group contains only one VD.
- OCE is not supported on RAID50 and RAID60. RLM is not supported on RAID10, RAID50 and RAID60.
- If the controller already contains the maximum number of virtual disks, you cannot perform a RAID level migration or capacity expansion on any virtual disk.
- The controller changes the write cache policy of all virtual disks undergoing a RLM/OCE to Write-Through until RLM/OCE is complete.
- Reconfiguring Virtual Disks typically impacts disk performance until the reconfiguration operation is complete.
- The total number of physical disks in a disk group cannot exceed 32.
- If any background operation (like BGI/rebuild/copyback/patrol read) is already running on the corresponding VD/PD then Reconfiguration (OCE/RLM) is not allowed at that time.
- Any kind of disk migration when Reconfiguration (OCE/RLM) is on progress on drives associated with VD causes reconfiguration to fail.
- Any new drive added for OCE/RLM becomes part of the VD after reconstruction completes. But State for those new drive changes to Online just after reconstruction starts.

Cancel Initialization

This feature is the ability to cancel the background initialization on a virtual disk. On PERC controllers, the background initialization of redundant virtual disk starts automatically after a virtual disk is created. The background initialization of redundant virtual disk prepares the virtual disk for parity information and improves write performance. However, some processes such as creating a virtual disk cannot be run while the background initialization is in progress. Cancel Initialization provides the ability to cancel the background initialization manually. Once cancelled, the background initialization automatically restarts within 0 to 5 minutes.

(i) NOTE: Background initialization is not applicable for RAID 0 virtual disks.

Managing virtual disks using web interface

- 1. In the iDRAC web interface, go to Configuration > Storage Configuration > Virtual Disk Configuration.
- 2. From the Virtual Disks, select the controller for which you want to manage the virtual disks.
- 3. From Action drop-down menu, select an action.
 - When you select an action, an additional Action window displayed. Select / enter the desired value.
 - Rename
 - Delete



- Edit Cache Policy You can change the cache policy for the following options:
 - **Read Policy** Following values are available for selection:
 - Adaptive Read Ahead Indicates that for the given volume, the control uses the Read-Ahead cache policy if the two most recent disks accesses occurred in sequential sectors. If the read requests are random, the controller returns to No Read Ahead mode.
 - No Read Ahead Indicates that for the given volume, no read ahead policy is used.
 - Read Ahead Indicates that for the given volume, the controller reads sequentially ahead of the requested data and stores the additional data in cache memory, anticipating a data requirement. This speeds up sequential data reads, but there is less improvement when accessing random data.
 - Write Policy Change the write cache policy to one of the following options:
 - Write Through Indicates that for the given volume, the controller sends a data transfer completion signal to the host system when the disk subsystem has received all the data in a transaction.
 - Write Back Indicates that for the given volume, the controller sends a data transfer completion signal to the host system when the controller cache has received all the data in a transaction. The controller then writes the cached data to the storage device in the background.
 - Force Write Back When using force write-back caching, the write cache is enabled regardless of whether the controller has a battery. If the controller does not have a battery and force write-back caching is used, data loss may occur in the event of a power failure.
 - **Disk Cache Policy** Change the disk cache policy to one of the following options:
 - **Default** Indicates that the disk is using its default write cache mode. For SATA disks, this is enabled and for SAS disks this is disabled.
 - **Enabled** Indicates that the disk's write cache is enabled. This increases performance and the probability of data loss if there is power loss.
 - Disabled Indicates that the disk's write cache is disabled. This decreases performance and the probability of data loss.
- Edit Disk Capacity You can add the physical disks to the selected virtual disk in this window. This window also shows the current capacity and new capacity of the virtual disk after adding the physical disks.
- **RAID Level Migration** Displays the Disk Name, Current RAID Level, and size of the virtual disk. Allows you to select a New RAID Level. User may have to add additional drives to existing Virtual disks to migrate to new raid level. This feature is not applicable on RAID 10, 50 and 60.
- Initialize: Fast Updates the metadata on the physical disks so that all the disk space is available for future write operations. The initialize option can be completed quickly because existing information on the physical disks is not erased, although future write operations overwrites any information that remains on the physical disks.
- Initialize: Full All existing data and file systems are erased.

(i) NOTE: The Initialize: Full option is not applicable for PERC H330 controllers.

• Check Consistency — To check the consistency of a virtual disk, select Check Consistency from the corresponding drop-down menu.

(i) NOTE: Consistency check is not supported on drives set up in RAID0 mode.

For more information about these options, see the *iDRAC Online Help*.

 Click Apply Now to apply the changes immediately, At Next Reboot to apply the changes after next reboot, At Scheduled Time to apply the changes at a particular time, and Discard All Pending to discard the changes. Based on the selected operation mode, the settings are applied.

Managing virtual disks using RACADM

Use the following commands to manage virtual disks:

• To delete virtual disk:

racadm storage deletevd:<VD FQDD>

To initialize virtual disk:

racadm storage init:<VD FQDD> -speed {fast|full}

To check consistency of virtual disks (not supported on RAID0):

```
racadm storage ccheck: <vdisk fqdd>
```

To cancel the consistency check:



racadm storage cancelcheck: <vdisks fqdd>

• To encrypt virtual disks:

racadm storage encryptvd:<VD FQDD>

• To assign or unassign dedicated hot spares:

```
racadm storage hotspare:<Physical Disk FQDD> -assign <option> -type dhs -vdkey: <FQDD
of VD>
```

<option>=yes

Assign hot spare

<Option>=NO

Unassign hot spare

RAID Configuration Features

Following table lists some of the RAID configuration features which are available in RACADM and WSMan:

CAUTION: Forcing a physical disk to go online or offline may result in data loss.

Table 56. RAID Configuration Features

Feature	RACADM Command	Description
Force Online	racadm storage forceonline: <pd fqdd=""></pd>	A power failure, corrupted data, or some other reason may lead to a physical disk going offline. You can use this feature to force a physical disk back into an online state when all other options have been exhausted. Once the command is run, the controller places the drive back into online state and restore its membership within the virtual disk. This happens only if the controller can read from the drive and can write into its metadata.
NOTE: Data recovery is only possible failed disk.	if a limited portion of the disk is damaged.	Force Online feature cannot fix an already
Force Offline	racadm storage forceoffline: <pd fqdd=""></pd>	This feature removes a drive from a virtual disk configuration so that it goes offline, resulting in a degraded VD configuration. It is helpful if a drive is likely to fail in near future or is reporting a SMART failure but is still online. It can be also used if you would like to utilize a drive which is part of an existing RAID configuration.
Replace Physical Disk	racadm storage replacephysicaldisk: <source PD FQDD > -dstpd <destination fqdd="" pd=""></destination></source 	Allows you to copy data from a physical disk which is a member of a VD, to another physical disk. The source disk should be in online state, while the destination disk should be in ready state and of a similar size and type to replace the source.
Virtual Disk as boot device	racadm storage setbootvd: <controller< td=""><td>A virtual disk can be configured as a boot device using this feature. This enables fault tolerance when a VD with</td></controller<>	A virtual disk can be configured as a boot device using this feature. This enables fault tolerance when a VD with



Table 56. RAID Configuration Features (continued)

Feature	RACADM Command	Description
	FQDD> -vd <virtualdisk FQDD></virtualdisk 	redundancy is selected as the boot device, and also has the operating system installed on it.
Unlock Foreign Configuration	racadm storage unlock: <controller fqdd=""> -key <key id=""> -passwd <passphrase></passphrase></key></controller>	This feature is used to authenticate locked drives which have a different source controller encryption than the destination. Once unlocked, the drive can be successfully migrated from one controller to another.

Managing controllers

You can perform the following for controllers:

- Configure controller properties
- Import or auto import foreign configuration
- Clear foreign configuration
- Reset controller configuration
- Create, change, or delete security keys
- Discard preserved cache

Configuring controller properties

You can configure the following properties for the controller:

- Patrol read mode (auto or manual)
- Start or stop patrol read if patrol read mode is manual
- Patrol read unconfigured areas
- Check consistency mode
- Copyback mode
- Load balance mode
- Check consistency rate
- Rebuild rate
- BGI rate
- Reconstruct rate
- Enhanced auto import foreign configuration
- Create or change security keys
- Encryption mode (Local Key Management and Secure Enterprise key Manager)

You must have Login and Server Control privilege to configure the controller properties.

Patrol read mode considerations

Patrol read identifies disk errors to avoid disk failures, data loss, or corruption. It runs automatically once a week on SAS and SATA HDDs.

The Patrol Read does not run on a physical disk in the following circumstances:

- The physical disk is an SSD.
- The physical disk is not included in a virtual disk or assigned as a hot spare.



- The physical disk is included in a virtual disk that is undergoing one of the following:
 - A rebuild
 - A reconfiguration or reconstruction
 - A background initialization
 - A check consistency

In addition, the Patrol Read operation suspends during heavy I/O activity and resumes when the I/O is complete.

NOTE: For more information on how often the Patrol Read operation runs when in auto mode, see the respective controller documentation.

() NOTE: Patrol read mode operations such as **Start** and **Stop** are not supported if there are no virtual disks available in the controller. Though you can invoke the operations successfully using the iDRAC interfaces, the operations fail when the associated job is started.

Load balance

The Load Balance property provides the ability to automatically use both controller ports or connectors connected to the same enclosure to route I/O requests. This property is available only on SAS controllers.

Bgi rate

(i) NOTE: Both H330 and H345 require the driver to be loaded for the background initialization operations to run.

On PERC controllers, background initialization of a redundant virtual disk begins automatically within 0 to 5 minutes after the virtual disk is created. The background initialization of a redundant virtual disk prepares the virtual disk to maintain redundant data and improves write performance. For example, after the background initialization of a RAID 5 virtual disk completes, the parity information has been initialized. After the background initialization of a RAID 1 virtual disk completes, the physical disks are mirrored.

The background initialization process helps the controller identify and correct problems that may occur with the redundant data later. In this regard, the background initialization process is similar to a check consistency. The background initialization should be allowed to run to completion. If cancelled, the background initialization automatically restarts within 0 to 5 minutes. Some processes such as read and write operations are possible while the background initialization is running. Other processes, such as creating a virtual disk, cannot be run concurrently with a background initialization. These processes cause the background initialization to cancel.

The background initialization rate, configurable between 0% and 100%, represents the percentage of the system resources dedicated to running the background initialization task. At 0%, the background initialization has the lowest priority for the controller, takes the most time to complete, and is the setting with the least impact to system performance. A background initialization rate of 0% does not mean that the background initialization is stopped or paused. At 100%, the background initialization is the highest priority for the controller. The background initialization time is minimized and is the setting with the most impact to system performance.

Check consistency

The Check Consistency task verifies the accuracy of the redundant (parity) information. This task only applies to redundant virtual disks. When necessary, the Check Consistency task rebuilds the redundant data. If the virtual disk is in a Failed Redundancy state, running a check consistency may be able to return the virtual disk to a Ready state.

The check consistency rate, configurable between 0% and 100%, represents the percentage of the system resources dedicated to running the check consistency task. At 0%, the check consistency has the lowest priority for the controller, takes the most time to complete, and is the setting with the least impact to system performance. A check consistency rate of 0% does not mean that the check consistency is stopped or paused. At 100%, the check consistency is the highest priority for the controller. The check consistency time is minimized and is the setting with the most impact to system performance.

Create or change security keys

When configuring the controller properties, you can create or change the security keys. The controller uses the encryption key to lock or unlock access to SED. You can create only one encryption key for each encryption-capable controller. The security key is managed using following features:



- Local Key Management (LKM) System LKM is used to generate the key ID and the password or key required to secure the virtual disk. If you are using LKM, you must create the encryption key by providing the Security Key Identifier and the Passphrase.
- 2. Secure Enterprise Key Manager (SEKM) This feature is used to generate the key using the Key Management Server (KMS). If you are using SEKM, you must configure iDRAC with KMS information as well as SSL related configuration.

(i) NOTE:

- This task is not supported on PERC hardware controllers running in eHBA mode.
- If you create the security key in 'Add to Pending Operation' mode and a job is not created, and then if you delete the security key, the create security key pending operation is cleared.

() NOTE:

- For enabling SEKM, ensure that the supported PERC firmware is installed.
- You cannot downgrade the PERC firmware to the previous version if SEKM is enabled. Downgrading of other PERC controller firmware in the same system which is not in SEKM mode may also fail. To downgrade the firmware for the PERC controllers that are not in SEKM mode, you can use OS DUP update method, or disable SEKM on the controllers and then retry the downgrade from iDRAC.

NOTE: When importing a hot plugged locked volume from one server to another, you will see CTL entries for Controller attributes being applied in the LC Log.

Configuring controller properties using web interface

- 1. In the iDRAC web interface, go to **Storage** > **Overview** > **Controllers**. The **Setup Controllers** page is displayed.
- 2. In the Controller section, select the controller that you want to configure.
- **3.** Specify the required information for the various properties.

The **Current Value** column displays the existing values for each property. You can modify this value by selecting the option from the **Action** drop-down menu for each property.

For information about the fields, see the *iDRAC Online Help*.

- 4. From the Apply Operation Mode, select when you want to apply the settings.
- 5. Click Apply.

Based on the selected operation mode, the settings are applied.

Configuring controller properties using RACADM

• To set Patrol Read Mode:

racadm set storage.controller.<index>.PatrolReadMode {Automatic | Manual | Disabled}

• If Patrol read mode is set to manual, use the following commands to start and stop Patrol read Mode:

racadm storage patrolread:<Controller FQDD> -state {start|stop}

- () NOTE: Patrol read mode operations such as Start and Stop are not supported if there are no virtual disks available in the controller. Though you can invoke the operations successfully using the iDRAC interfaces, the operations will fail when the associated job is started.
- To specify the Check Consistency Mode, use Storage.Controller.CheckConsistencyMode object.
- To enable or disable the Copyback Mode, use Storage.Controller.CopybackMode object.
- To enable or disable the Load Balance Mode, use Storage.Controller.PossibleloadBalancedMode object.
- To specify the percentage of the system's resources dedicated to perform a check consistency on a redundant virtual disk, use **Storage.Controller.CheckConsistencyRate** object.
- To specify the percentage of the controller's resources dedicated to rebuild a failed disk, use Storage.Controller.RebuildRate object



- To specify the percentage of the controller's resources dedicated to perform the background initialization (BGI) of a virtual disk after it is created, use**Storage.Controller.BackgroundInitializationRate** object
- To specify the percentage of the controller's resources dedicated to reconstruct a disk group after adding a physical disk or changing the RAID level of a virtual disk residing on the disk group, use **Storage.Controller.ReconstructRate** object
- To enable or disable the enhanced auto import of foreign configuration for the controller, use **Storage.Controller.EnhancedAutoImportForeignConfig** object
- To create, modify, or delete security key to encrypt virtual drives:

```
racadm storage createsecuritykey:<Controller FQDD> -key <Key id> -passwd <passphrase>
racadm storage modifysecuritykey:<Controller FQDD> -key <key id> -oldpasswd <old
passphrase> -newpasswd <new passphrase>
racadm storage deletesecuritykey:<Controller FQDD>
```

Importing or auto importing foreign configuration

A foreign configuration is data residing on physical disks that have been moved from one controller to another. Virtual disks residing on physical disks that have been moved are considered to be a foreign configuration.

You can import foreign configurations so that virtual disks are not lost after moving Physical Disks. A foreign configuration can be imported only if it contains a virtual disk that is in either Ready or Degraded state or a hotspare that is dedicated to a virtual disk which can be imported or is already present.

All of the virtual disk data must be present, but if the virtual disk is using a redundant RAID level, the additional redundant data is not required.

For example, if the foreign configuration contains only one side of a mirror in a RAID 1 virtual disk, then the virtual disk is in a Degraded state and can be imported. If the foreign configuration contains only one physical disk that was originally configured as a RAID 5 using three physical disks, then the RAID 5 virtual disk is in a Failed state and cannot be imported.

In addition to virtual disks, a foreign configuration may consist of a physical disk that was assigned as a hot spare on one controller and then moved to another controller. The Import Foreign Configuration task imports the new physical disk as a hot spare. If the physical disk was set as a dedicated hot spare on the previous controller, but the virtual disk to which the hot spare was assigned is no longer present in the foreign configuration, then the physical disk is imported as a global hot spare.

If any foreign configurations locked using Local Key manager (LKM) are Detected, then import foreign configuration operation is not possible in iDRAC in this release. You must unlock the drives through CTRL-R and then continue to import foreign configuration from iDRAC.

The Import Foreign Configuration task is only displayed when the controller has detected a foreign configuration. You can also identify whether a physical disk contains a foreign configuration (virtual disk or hot spare) by checking the physical disk state. If the physical disk state is Foreign, then the physical disk contains all or some portion of a virtual disk or has a hot spare assignment.

() NOTE: The task of importing foreign configuration imports all virtual disks residing on physical disks that have been added to the controller. If more than one foreign virtual disk is present, all the configurations are imported.

PERC9 controller provides support for auto import of foreign configuration without requiring user interactions. The auto import can be enabled or disabled. If enabled, the PERC controller can auto import any foreign configuration detected without manual intervention. If disabled the PERC does not auto import any foreign configuration.

You must have Login and Server Control privilege to import foreign configurations.

This task is not supported on PERC hardware controllers running in HBA mode.

NOTE: It is not recommended to remove an external enclosure cable while the operating system is running on the system. Removing the cable could result in a foreign configuration when the connection is re-established.

You can manage foreign configurations in the following cases:

- All the physical disks in a configuration are removed and re-inserted.
- Some of the physical disks in a configuration are removed and re-inserted.
- All the physical disks in a virtual disk are removed, but at different times, and then re-inserted.
- The physical disks in a non-redundant virtual disk are removed.

The following constraints apply to the physical disks that are considered for import:



- The drive state of a physical disk can change from the time the foreign configuration is scanned to when the actual import
 occurs. The foreign import occurs only on drives that are in the Unconfigured Good state.
- Drives in the failed or offline state cannot be imported.
- The firmware does not allow you to import more than eight foreign configurations.

Importing foreign configuration using web interface

NOTE: If there is an incomplete foreign disk configuration in the system, then the state of one or more existing online virtual disks is also displayed as foreign.

(i) NOTE: Importing foreign configuration for BOSS controller is not supported.

To import foreign configuration:

- 1. In the iDRAC web interface, go to Configuration > Storage Configuration;.
- 2. From the Controller drop-down menu, select the controller you want to import the foreign configuration to.
- 3. Click Import under the Foreign Configuration and then click Apply.

Importing foreign configuration using RACADM

To import foreign configuration:

racadm storage importconfig:<Controller FQDD>

For more information, see the iDRAC RACADM Command Line Reference Guide available at dell.com/idracmanuals.

Clearing foreign configuration

After moving a physical disk from one controller to another, you may find that the physical disk contains all or some portion of a virtual disk (foreign configuration). You can identify whether a previously used physical disk contains a foreign configuration (virtual disk) by checking the physical disk state. If the physical disk state is Foreign, then the physical disk contains all or some portion of a virtual disk. You can clear or erase the virtual disk information from the newly attached physical disks.

The Clear Foreign Configuration operation permanently erases all data residing on the physical disks that are added to the controller. If more than one foreign virtual disk is present, all the configurations are erased. You may prefer to import the virtual disk rather than destroy the data. An initialization must be performed to remove foreign data. If you have an incomplete foreign configuration which cannot be imported, you can use the Clearing Foreign Configuration option to erase the foreign data on the physical disks.

Clearing foreign configuration using web interface

To clear the foreign configuration:

- 1. In the iDRAC web interface, go to Configuration > Storage Configuration > Controller Configuration. The Controller Configuration page is displayed.
- 2. From the **Controller** drop-down menu, select the controller for which you want to clear the foreign configuration.

(i) NOTE: To clear foreign configuration on BOSS controllers, click "Reset Configuration".

3. Click Clear Configuration.

4. Click Apply

Based on the selected operation mode, the virtual disks residing on the physical disk is erased.



Clearing foreign configuration using RACADM

To clear foreign configuration:

racadm storage clearconfig:<Controller FQDD>

For more information, see the *iDRAC RACADM Command Line Reference Guide* available at **dell.com/idracmanuals**.

Resetting controller configuration

You can reset the configuration for a controller. This operation deletes virtual disk drives and unassigns all hot spares on the controller. It does not erase any data other than removing the disks from the configuration. Reset configuration also does not remove any foreign configurations. The real-time support of this feature is available only in PERC 9.1 firmware. Reset configuration does not erase any data. You may recreate the exact same configuration without an initialize operation which may result in the data being recovered. You must have server control privilege.

NOTE: Resetting the controller configuration does not remove a foreign configuration. To remove a foreign configuration, perform clear configuration operation.

Resetting controller configuration using web interface

To reset the controller configuration:

- 1. In the iDRAC Web interface, go to **Storage** > **Overview** > **Controllers**.
- 2. From the Actions, select Reset Configuration for one or more controllers.
- 3. For each controller, from the Apply Operation Mode drop-down menu, select when you want to apply the settings.
- Click Apply.
 Based on the selected operation mode, the settings are applied.

Resetting controller configuration using RACADM

To reset the controller configuration:

racadm storage resetconfig:<Controller FQDD>

For more information, see the iDRAC RACADM Command Line Reference Guide available at dell.com/idracmanuals.

Switching the controller mode

On PERC 9.1 controllers, you can change the personality of the controller by switching the mode from RAID to HBA. The controller operates similar to an HBA controller where the drivers are passed through the operating system. The controller mode change is a staged operation and does not occur in real time.

PERC 10 and later controllers supports Enhanced HBA mode, replacing HBA from the current controller mode options. However, PERC 9 still continues to support the HBA mode.

() NOTE:

- Enhanced HBA supports Non-RAID PDs and all RAID level VDs.
- It only supports creation of RAID0, RAID1, and RAID10 VDs.
- Enhanced HBA is not supported on PERC 11.

Enhanced HBA mode provides the following features:

- Create virtual disks with RAID level 0, 1, or 10.
- Present non-RAID disks to host.
- Configure a default cache policy for virtual disks as write-back with read ahead.
- Configure virtual disks and non-RAID disks as valid boot devices.
- Automatically convert all unconfigured disks to non-RAID:
 - On system boot



- On controller reset
- When unconfigured disks are hot-inserted

NOTE: Creating or importing RAID 5, 6, 50, or 60 virtual disks is not supported. Also, in enhanced HBA mode, non-RAID disks are enumerated first in ascending order, while RAID volumes are enumerated in descending order.

Before you change the mode of the controller from RAID to HBA, ensure that:

- The RAID controller supports the controller mode change. The option to change the controller mode is not available on controllers where the RAID personality requires a license.
- All virtual disks must be deleted or removed.
- Hot spares must be deleted or removed.
- Foreign configurations must be deleted or cleared.
- All physical disks that are in a failed state, must be removed or the pinned cache needs to be cleared.
- Any local security key that is associated with SEDs must be deleted.
- The controller must not have preserved cache.
- You have server control privileges to switch the controller mode.
- () NOTE: Ensure that you back up the foreign configuration, security key, virtual disks, and hot spares before you switch the mode as the data is deleted.
- () NOTE: Ensure that a CMC license (not applicable for MX platforms) is available for PERC FD33xS and FD33xD storage sleds before you change the controller mode. For more information on CMC license for the storage sleds, see the *Dell Chassis Management Controller Version 1.2 for PowerEdge FX2/FX2s User's Guide* available at dell.com/cmcmanuals.

Exceptions while switching the controller mode

The following list provides the exceptions while setting the controller mode using the iDRAC interfaces such as web interface, RACADM, and WSMan:

- If the PERC controller is in RAID mode, you must clear any virtual disks, hot spares, foreign configurations, controller keys, or
 preserved cache before changing it to HBA mode.
- You cannot configure other RAID operations while setting the controller mode. For example, if the PERC is in RAID mode and you set the pending value of the PERC to HBA mode, and you try to set the BGI attribute, the pending value is not initiated.
- When you switch the PERC controller from HBA to RAID mode, the drives remain in Non-RAID state and are not automatically set to Ready state. Additionally, the **RAIDEnhancedAutoImportForeignConfig** attribute is automatically set to **Enabled**.

The following list provides the exceptions while setting the controller mode using the Server Configuration Profile feature using the WSMan or RACADM interface:

- Server Configuration Profile feature allows you to configure multiple RAID operations along with setting the controller mode. For example, if the PERC controller is in HBA mode, you can edit the export Server Configuration Profile (SCP) to change the controller mode to RAID, convert drives to ready and create a virtual disk.
- While changing the mode from RAID to HBA, the **RAIDaction pseudo** attribute is set to update (default behavior). The attribute runs and creates a virtual disk which fails. The controller mode is changed, however, the job is completed with errors. To avoid this issue, you must comment out the RAIDaction attribute in the SCP file.
- When the PERC controller is in HBA mode, if you run import preview on export SCP which is edited to change controller mode to RAID, and try creating a VD, the virtual disk creation fails. Import preview does not support validating stacking RAID operations with changing controller mode.

Switching the controller mode using the iDRAC web interface

To switch the controller mode, perform the following steps:

- 1. In the iDRAC web interface, click Storage > Overview > Controllers.
- 2. On the Controllers page, click Action > Edit.The Current Value column displays the current setting of the controller.
- **3.** From the drop-down menu, select the controller mode you want to switch to, and click **At Next Reboot**. Reboot the system for the change to take effect.



Switching the controller mode using RACADM

To switch the controller mode using RACADM, run the following commands.

To view the current mode of the controller:

\$ racadm get Storage.Controller.1.RequestedControllerMode[key=<Controller_FQDD>]

The following output is displayed:

RequestedControllerMode = NONE

• To set the controller mode as HBA:

\$ racadm set Storage.Controller.1.RequestedControllerMode HBA [Key=<Controller FQDD>]

• To create a job and apply changes:

\$ racadm jobqueue create <Controller Instance ID> -s TIME_NOW -r pwrcycle

For more information, see the *iDRAC RACADM Command Line Interface Reference Guide* available at **dell.com/idracmanuals**.

12 Gbps SAS HBA adapter operations

Dell PowerEdge servers must have an operating system installed and the appropriate device driver to be loaded in order for Dell HBAs to operate. Following POST, the HBA ports will be disabled. The HBA device driver is responsible for resetting the HBA and enabling its ports connected to storage devices. Without an operating system, the driver will not be loaded, and there is no guarantee that iDRAC will be able to display storage devices connected to Dell HBAs.

The non-RAID controllers are the HBAs that do not have few RAID capabilities. They do not support virtual disks.

14G iDRAC interface supports 12 Gbps SAS HBA controller, HBA330 (integrated and adapter) controllers, HBA330 MMZ, and HBA330 MX adapters.

AMD platforms support HBA355i front and HBA355i Adapter controllers.

You can perform the following for non-RAID controllers:

- View controller, physical disks, and enclosure properties as applicable for the non-RAID controller. Also, view EMM, fan, power supply unit, and temperature probe properties associated with the enclosure. The properties are displayed based on the type of controller.
- View software and hardware inventory information.
- Update firmware for enclosures behind the 12 Gbps SAS HBA controller (staged)
- Monitor the polling or polling frequency for physical disk SMART trip status when there is change detected
- Monitor the physical disks hot plug or hot removal status
- Blink or unblink LEDs

() NOTE:

- There is limited support for tape drives when they are connected behind 12gbps SAS or HBA355e.
- Even though LED is not available for tape drive, blink/unblink option can be successful.

() NOTE:

- Enable Collect System Inventory On Reboot (CSIOR) operation before inventorying or monitoring the non-RAID controllers.
- Real-time monitoring for SMART enabled drives and SES enclosure sensors is only done for the 12 Gbps SAS HBA controllers and HBA330 internal controllers.

(i) NOTE: Detection of failed drives behind SAS HBA controllers is not supported.

Monitoring predictive failure analysis on drives

Storage management supports Self Monitoring Analysis and Reporting Technology (SMART) on physical disks that are SMARTenabled.



SMART performs predictive failure analysis on each disk and sends alerts if a disk failure is predicted. The controllers check physical disks for failure predictions and, if found, pass this information to iDRAC. iDRAC immediately logs an alert.

Controller operations in non-RAID mode or HBA mode

If the controller is in non-RAID mode (HBA mode), then:

- Virtual disks or hot spares are not available.
- Security state of the controller is disabled.
- All physical disks are in non-RAID mode.

You can perform the following operations if the controller is in non-RAID mode:

- Blink/unblink the physical disk.
- Configure all properties including the following:
- Load balanced mode
- Check consistency mode
- Patrol read mode
- Copyback mode
- Controller boot mode
- Enhanced auto import foreign configuration
- Rebuild rate
- Check consistency rate
- Reconstruct rate
- o BGI rate
- Enclosure or backplane mode
- Patrol read unconfigured areas
- View all properties that are applicable to a RAID controller expect for virtual disks.
- Clear foreign configuration

(i) NOTE: If an operation is not supported in non-RAID mode, an error message is displayed.

You cannot monitor the enclosure temperature probes, fans, and power supplies when the controller is in non-RAID mode.

Running RAID configuration jobs on multiple storage controllers

While performing operations on more than two storage controllers from any supported iDRAC interface, make sure to:

- Run the jobs on each controller individually. Wait for each job to complete before starting the configuration and job creation on the next controller.
- Schedule multiple jobs to run at a later time using the scheduling options.

Manage Preserved cache

The Managed Preserved Cache feature is a controller option which provides the user an option to discard the controller cache data. In the write-back policy, data is written to the cache before being written to the physical disk. If the virtual disk goes offline or is deleted for any reason, the data in the cache gets deleted.

The PREC Controller preserves the data written on the preserved or dirty cache in an event of power failure or cable disconnect until you recover the virtual disk or clear the cache.

The status of the controller is affected by the preserved cache. The controller status is displayed as degraded if the controller has preserved cache. Discard the preserved cache is possible only if all of the following conditions are met:

- The controller does not have any foreign configuration.
- The controller does not have any offline or missing virtual disks.
- Cables to any virtual disk are not disconnected.



Managing PCIe SSDs

Peripheral Component Interconnect Express (PCIe) solid-state device (SSD) is a high-performance storage device designed for solutions requiring low latency, high Input Output Operations per Second (IOPS), and enterprise class storage reliability and serviceability. The PCIe SSD is designed based on Single Level Cell (SLC) and Multi-Level Cell (MLC) NAND flash technology with a high-speed PCIe 2.0, PCIe 3.0, or PCIe 4.0 compliant interface. In 14th generation of PowerEdge servers, we have three different ways to connect SSDs. You can use an extender to connect the SSDs via backplane, directly connect the SSDs from backplane to mother board using slimline cable without extender, and use HHHL (Add-In) card which sits on the motherboard.

() NOTE:

- 14th generation of PowerEdge servers are supporting Industry standard NVMe-MI specification based NVMe SSDs
- PERC 11 supports PCIe SSD/NVMe devices behind PERC inventory monitoring and configuration.

Using iDRAC interfaces, you can view and configure NVMe PCIe SSDs.

The key features of PCIe SSD are:

- Hot plug capability
- High-performance device

In few of the 14th generation of PowerEdge servers, up to 32 NVMe SSDs are supported.

You can perform the following operations for PCIe SSDs:

- Inventory and remotely monitor the health of PCIe SSDs in the server
- Prepare to remove the PCIe SSD
- Securely erase the data
- Blink or unblink the device LED (Identify the device)
- You can perform the following operations for HHHL SSDs:
- Inventory and real-time monitoring of the HHHL SSD in the server
- Failed card reporting and logging in iDRAC and OMSS
- Securely erasing the data and removing the card
- TTY logs reporting
- You can perform the following operations for SSDs:
- Drive status reporting such as Online, Failed, and Offline
- **NOTE:** Hot plug capability, prepare to remove, and blink or unblink the device LED is not applicable for HHHL PCIe SSD devices.
- **NOTE:** When NVMe devices are controlled behind S140, prepare to remove and cryptographic erase operations are not supported, blink and unblink are supported.

Inventorying and monitoring PCIe SSDs

The following inventory and monitoring information is available for PCle SSDs:

- Hardware information:
- PCIe SSD Extender card
- PCIe SSD Backplane

If the system has a dedicated PCIe backplane, two FQDDs are displayed. One FQDD is for regular drives and the other is for SSDs. If the backplane is shared (universal), only one FQDD is displayed. In case the SSDs are directly attached, the controller FQDD reports as CPU.1, indicating that the SSD is directly attached to the CPU.

• Software inventory includes only the firmware version for the PCIe SSD.

Inventorying and monitoring PCIe SSDs using web interface

To inventory and monitor PCle SSD devices, in the iDRAC web interface, go to **Storage** > **Overview** > **Physical Disks**. The **Properties** page is displayed. For PCle SSDs, the **Name** column displays **PCle SSD**. Expand to view the properties.



Inventorying and monitoring PCIe SSDs using RACADM

Use the racadm storage get controllers: <PcieSSD controller FQDD> command to inventory and monitor PCle SSDs.

To view all PCIe SSD drives:

racadm storage get pdisks

To view PCIe extender cards:

racadm storage get controllers

To view PCIe SSD backplane information:

racadm storage get enclosures

(i) NOTE: For all the mentioned commands, PERC devices are also displayed.

For more information, see the iDRAC RACADM Command Line Reference Guide available at dell.com/idracmanuals.

Preparing to remove PCIe SSD

() NOTE: This operation is not supported when:

- PCIe SSD is configured using the S140 controller.
- NVMe device is behind PERC 11.

PCIe SSDs support orderly hot swap allowing you to add or remove a device without halting or rebooting the system in which the devices are installed. To prevent data loss, you must use the Prepare to Remove operation before physically removing a device.

Orderly hot swap is supported only when PCIe SSDs are installed in a supported system running a supported operating system. To ensure that you have the correct configuration for your PCIe SSD, see the system-specific owner's manual.

The Prepare to Remove operation is not supported for PCIe SSDs on the VMware vSphere (ESXi) systems and HHHL PCIe SSD devices.

NOTE: Prepare to Remove operation is supported on systems with ESXi 6.0 with iDRAC Service Module version 2.1 or higher.

The Prepare to Remove operation can be performed in real-time using iDRAC Service Module.

The Prepare to Remove operation stops any background activity and any ongoing I/O activity so that device can be removed safely. It causes the status LEDs on the device to blink. You can safely remove the device from the system under the following conditions after you initiate the Prepare to Remove operation:

- The PCIe SSD is blinking the safe to remove LED pattern (blinks amber).
- The PCIe SSD is no longer accessible by the system.

Before preparing the PCle SSD for removal, ensure that:

- iDRAC Service Module is installed.
- Lifecycle Controller is enabled.
- You have Server Control and Login privileges.

Preparing to remove PCIe SSD using web interface

To prepare the PCIe SSD for removal:

- In the iDRAC Web interface, go to Storage > Overview > Physical Disks. The Setup Physical Disk page is displayed.
- 2. From the Controller drop-down menu, select the extender to view the associated PCIe SSDs.
- 3. From the drop-down menus, select **Prepare to Remove** for one or more PCIe SSDs.

If you have selected **Prepare to Remove** and you want to view the other options in the drop-down menu, then select **Action** and then click the drop-down menu to view the other options.



(i) NOTE: Ensure that iSM is installed and running to perform the preparetoremove operation.

From the Apply Operation Mode drop-down menu, select Apply Now to apply the actions immediately.
 If there are jobs to be completed, then this option is grayed-out.

(i) NOTE: For PCIe SSD devices, only the Apply Now option is available. This operation is not supported in staged mode.

5. Click Apply.

If the job is not created, a message indicating that the job creation was not successful is displayed. Also, the message ID and the recommended response action is displayed.

If the job is created successfully, a message indicating that the job ID is created for the selected controller is displayed. Click **Job Queue** to view the progress of the job in the **Job Queue** page.

If pending operation is not created, an error message is displayed. If pending operation is successful and job creation is not successful, then an error message is displayed.

Preparing to remove PCIe SSD using RACADM

To prepare the PCIeSSD drive for removal:

racadm storage preparetoremove:<PCIeSSD FQDD>

To create the target job after executing preparetoremove command:

racadm jobqueue create <PCIe SSD FQDD> -s TIME_NOW --realtime

To query the job ID returned:

racadm jobqueue view -i <job ID>

For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Erasing PCIe SSD device data

(i) NOTE: This operation is not supported when PCIe SSD is configured using the SWRAID controller.

Cryptographic Erase permanently erases all data present on the disk. Performing a Cryptographic Erase on an PCIe SSD overwrites all blocks and results in permanent loss of all data on the PCIe SSD. During Cryptographic Erase, the host is unable to access the PCIe SSD. The changes are applied after system reboot.

If the system reboots or experiences a power loss during cryptographic erase, the operation is canceled. You must reboot the system and restart the process.

Before erasing PCIe SSD device data, make sure that:

- Lifecycle Controller is enabled.
- You have Server Control and Login privileges.

() NOTE:

- Erasing PCIe SSDs can only be performed as a staged operation.
- After the drive is erased, it displays in the operating system as online but it is not initialized. You must initialize and format the drive before using it again.
- After you hot-plug a PCIe SSD, it may take several seconds to appear on the web interface.

Erasing PCIe SSD device data using web interface

To erase the data on the PCIe SSD device:

 In the iDRAC Web interface, go to Storage > Overview > Physical Disks. The Physical Disk page is displayed.



- 2. From the **Controller** drop-down menu, select the controller to view the associated PCle SSDs.
- 3. From the drop-down menus, select **Cryptographic Erase** for one or more PCIe SSDs.

If you have selected **Cryptographic Erase** and you want to view the other options in the drop-down menu, then select **Action** and then click the drop-down menu to view the other options.

- 4. From the Apply Operation Mode drop-down menu, select one of the following options:
 - At Next Reboot Select this option to apply the actions during the next system reboot.
 - **At Scheduled Time** Select this option to apply the actions at a scheduled day and time:
 - Start Time and End Time Click the calendar icons and select the days. From the drop-down menus, select the time. The action is applied between the start time and end time.
 - From the drop-down menu, select the type of reboot:
 - No Reboot (Manually Reboot System)
 - Graceful Shutdown
 - Force Shutdown
 - Power Cycle System (cold boot)

5. Click Apply.

If the job is not created, a message indicating that the job creation was not successful is displayed. Also, the message ID and the recommended response action is displayed.

If the job is created successfully, a message indicating that the job ID is created for the selected controller is displayed. Click **Job Queue** to view the progress of the job in the Job Queue page.

If pending operation is not created, an error message is displayed. If pending operation is successful and job creation is not successful, then an error message is displayed.

Erasing PCIe SSD device data using RACADM

To securely erase a PCIe SSD device:

```
racadm storage secureerase:<PCIeSSD FQDD>
```

To create the target job after executing the secureerase command:

racadm jobqueue create <PCIe SSD FQDD> -s TIME NOW -e <start time>

To query the job ID returned:

racadm jobqueue view -i <job ID>

For more information, see the iDRAC RACADM Command Line Reference Guide available at dell.com/idracmanuals.

Managing enclosures or backplanes

You can perform the following for enclosures or backplanes:

- View properties
- Configure universal mode or split mode
- View slot information (universal or shared)
- Set SGPIO mode
- Set Asset Tag
- Asset Name

Configuring backplane mode

The Dell 14th generation PowerEdge servers supports a new internal storage topology, where two storage controllers (PERCs) can be connected to a set of internal drives through a single expander. This configuration is used for high performance mode with no failover or High Availability (HA) functionality. The expander splits the internal drive array between the two storage



controllers. In this mode, virtual disk creation only displays the drives connected to a particular controller. There are no licensing requirements for this feature. This feature is supported only on a few systems.

Backplane supports the following modes:

- Unified mode This is the default mode. The primary PERC controller has access to all the drives connected to the backplane even if a second PERC controller is installed.
- Split mode One controller has access to the first 12 drives and the second controller has access to the last 12 drives. The drives connected to the first controller are numbered 0-11 while the drives connected to the second controller are numbered 12-23.
- Split mode 4:20 One controller has access to the first 4 drives and the second controller has access to the last 20 drives. The drives connected to the first controller are numbered 0-3 while the drives connected to the second controller are numbered 4-23.
- Split mode 8:16 One controller has access to the first 8 drives and the second controller has access to the last 16 drives. The drives connected to the first controller are numbered 0-7 while the drives connected to the second controller are numbered 8-23.
- Split mode 16:8 One controller has access to the first 16 drives and the second controller has access to the last 8 drives. The drives connected to the first controller are numbered 0-15 while the drives connected to the second controller are numbered 16-23.
- Split mode 20:4 One controller has access to the first 20 drives and the second controller has access to the last 4 drives. The drives connected to the first controller are numbered 0-19 while the drives connected to the second controller are numbered 20-23.
- Split mode 6:6:6:6 4 blades installed in one chassis and each blade have 6 drives assigned. This mode is only supported on PowerEdge C series blades.
- Information Not Available Controller information is not available.

iDRAC allows the split mode setting if the expander has the capability to support the configuration. Ensure that you enable this mode prior to installing the second controller. iDRAC performs a check for expander capability prior to allowing this mode to be configured and does not check whether the second PERC controller is present.

() NOTE: Cable errors (or other errors) may appear if you put the backplane into split mode with only one PERC attached, or if you put the backplane into Unified Mode with two PERCs attached.

To modify the setting, you must have Server Control privilege.

If any other RAID operations are in pending state or any RAID job is scheduled, you cannot change the backplane mode. Similarly, if this setting is pending, you cannot schedule other RAID jobs.

() NOTE:

- Warning messages are displayed when the setting is being changed as there is a possibility of data loss.
- LC Wipe or iDRAC reset operations do not change the expander setting for this mode.
- This operation is supported only in real-time and not staged.
- You can change the backplane configuration multiple times.
- The backplane splitting operation can cause data loss or foreign configuration if the drive association changes from one controller to another controller.
- During the backplane splitting operation, the RAID configuration may be impacted depending on the drive association.

Any change in this setting only takes effect after a system power reset. If you change from Split mode to Unified, an error message is displayed on the next boot as the second controller does not see any drives. Also, the first controller will see a foreign configuration. If you ignore the error, the existing virtual disks are lost.

Configuring backplane mode using web interface

To configure backplane mode using iDRAC web interface:

- 1. In the iDRAC web interface, go to Configuration > Storage Configuration > Enclosures Configuration.
- 2. From the **Controller** menu, select the controller to configure its associated enclosures.
- 3. From the Action drop-down menu, select Edit Enclosure Mode.
- The Edit Enclosure Mode page is displayed.
- 4. In the **Current Value** column, select the required enclosure mode for the backplane or enclosure. The options are:
 - Unified Mode
 - Split Mode
 - Split Mode 4:20



- Split Mode 8:16
- Split Mode 16:8
- Split Mode 20:4

NOTE: For C6420, the available modes are: Split Mode and Split Mode-6:6:6:6. Few values may be only supported on certain platforms.

For R740xd and R940, power cycle of the server is needed to apply the new backplane zone and for C6420, A/C cycle (of the blade chassis) to apply the new backplane zone.

- 5. Click Add to Pending Operations. A job ID is created.
- 6. Click Apply Now.
- 7. Go to the **Job Queue** page and verify that it displays the status as Completed for the job.
- 8. Power cycle the system for the setting to take effect.

Configuring enclosure using RACADM

To configure the enclosure or backplane, use the set command with the objects in **BackplaneMode**.

For example, to set the BackplaneMode attribute to split mode:

1. Run the following command to view the current backplane mode:

racadm get storage.enclosure.1.backplanecurrentmode

The output is:

BackplaneCurrentMode=UnifiedMode

2. Run the following command to view the requested mode:

racadm get storage.enclosure.1.backplanerequestedmode

The output is:

BackplaneRequestedMode=None

3. Run the following command to set the requested backplane mode to split mode:

racadm set storage.enclosure.1.backplanerequestedmode "splitmode"

The message is displayed indicating that the command is successful.

4. Run the following command to verify if the **backplanerequestedmode** attribute is set to split mode:

racadm get storage.enclosure.1.backplanerequestedmode

The output is:

BackplaneRequestedMode=None (Pending=SplitMode)

- 5. Run storage get controllers command and note down the controller instance ID.
- 6. Run the following command to create a job:

racadm jobqueue create <controller instance ID> -s TIME_NOW --realtime

A job ID is returned.

7. Run the following command to query the job status:

racadm jobqueue view -i JID_xxxxxxx

264 Managing storage devices



where, JID XXXXXXX is the job ID from step 6.

The status is displayed as Pending.

Continue to query the job ID until you view the Completed status (this process may take up to three minutes).

 $\textbf{8.} \ \ \mbox{Run the following command to view the backplanerequested mode attribute value:}$

 $\verb| racadm get storage.enclosure.1.backplanerequestedmode| \\$

The output is:

BackplaneRequestedMode=SplitMode

9. Run the following command to cold reboot the server:

racadm serveraction powercycle

10. After the system completes POST and CSIOR, type the following command to verify the backplanerequestedmode:

racadm get storage.enclosure.1.backplanerequestedmode

The output is:

BackplaneRequestedMode=None

11. Run the following to verify is the backplane mode is set to split mode:

racadm get storage.enclosure.1.backplanecurrentmode

The output is:

BackplaneCurrentMode=SplitMode

12. Run the following command and verify that only 0–11 drives are displayed:

racadm storage get pdisks

For more information about the RACADM commands, see the *iDRAC RACADM Command Line Interface Reference Guide* available at **dell.com/idracmanuals**.

Viewing universal slots

Some 14th generation PowerEdge server backplanes supports both SAS/SATA and PCIe SSD drives in the same slot. These slots are called universal slots and are wired to the primary storage controller (PERC) and either PCIe extender card or direct connect manager by CPU backplanes supports both SAS/SATA and PCIe SSD drives in the same slot. The backplane firmware provides information about the slots that support this feature. The backplane supports SAS/SATA disks or PCIe SSDs. Typically, the four higher number slots are universal. For example, in a universal backplane that supports 24 slots, slots 0-19 support only SAS/SATA disks while slots 20-23 support either SAS/SATA or PCIe SSD.

The roll-up health status for the enclosure provides the combined health status for all the drives in the enclosure. The enclosure link on the **Topology** page displays the entire enclosure information irrespective of which controller it is associated with. Because two storage controllers (PERC and PCle extender) can be connected to the same backplane, only the backplane associated with the PERC controller is displayed in **System Inventory** page.

In the Storage > Enclosures > Properties page, the Physical Disks Overview section displays the following:

- **Slot Empty** If a slot is empty.
- PCIe Capable If there are no PCIe capable slots, this column is not displayed.
- Bus Protocol If it is a universal backplane with PCIe SSD installed in one of the slots, this column displays PCIe.
- Hotspare This column is not applicable for PCIe SSD.

() NOTE: Hot swapping is supported for universal slots. If you want to remove a PCIe SSD drive and swap it with a SAS/ SATA drive, ensure that you first complete the PrepareToRemove task for the PCIe SSD drive. If you do not perform this task, the host operating system may have issues such as a blue screen, kernel panic, and so on.



Setting SGPIO mode

The storage controller can connect to the backplane in I2C mode (default setting for Dell backplanes) or Serial General Purpose Input/Output (SGPIO) mode. This connection is required for blinking LEDs on the drives. Dell PERC controllers and backplane support both these modes. To support certain channel adapters, the backplane mode must be changed SGPIO mode.

The SGPIO mode is only supported for passive backplanes. It is not supported for expander-based backplanes or passive backplanes in downstream mode. Backplane firmware provides information on capability, current state, and requested state.

After LC wipe operation or iDRAC reset to default, the SGPIO mode is reset to disabled state. It compares the iDRAC setting with the backplane setting. If the backplane is set to SGPIO mode, iDRAC changes its setting to match the backplane setting.

Server power cycle is required for any change in setting to take effect.

You must have Server Control privilege to modify this setting.

(i) NOTE: You cannot set the SGPIO mode using iDRAC Web interface.

Setting SGPIO mode using RACADM

To configure the SGPIO mode, use the set command with the objects in the SGPIOMode group.

If it is set to disabled, it is I2C mode. If enabled, it is set to SGPIO mode.

For more information, see the *iDRAC RACADM Command Line Interface Reference Guide* available at **dell.com/idracmanuals**.

Set Enclosure Asset Tag

Set Enclosure Asset Tag allows you to configure Asset Tag of a storage enclosure.

User can change the Asset Tag property of the enclosure to identify enclosures. These fields are checked for invalid values and an error is displayed if an invalid value is entered. These fields are part of the enclosure firmware; the data initially shown are the values saved in the firmware.

(i) NOTE: Asset Tag has a character limit of 10 that includes the null character.

(i) NOTE: These operations are not supported on internal enclosures.

Set Enclosure Asset Name

Set Enclosure Asset Name allows the user to configure the Asset Name of a storage enclosure.

The user can change the Asset Name property of the enclosure to identify enclosures easily. These fields are checked for invalid values and an error is displayed if an invalid value is entered. These fields are part of the enclosure firmware; the data initially shown are the values saved in the firmware.

(i) NOTE: Asset Name has a character limit of 32 that includes the null character.

(i) NOTE: These operations are not supported on internal enclosures.

Choosing operation mode to apply settings

While creating and managing virtual disks, setting up physical disks, controllers, and enclosures or resetting controllers, before you apply the various settings, you must select the operation mode. That is, specify when you want to apply the settings:

- Immediately
- During the next system reboot
- At a scheduled time
- As a pending operation to be applied as a batch as part of a single job.



Choosing operation mode using web interface

To select the operation mode to apply the settings:

- 1. You can select the operation mode on when you are on any of the following pages:
 - Storage > Physical Disks .
 - Storage > Virtual Disks
 - Storage > Controllers
 - Storage > Enclosures
- 2. Select one of the following from the Apply Operation Mode drop-down menu:
 - **Apply Now** Select this option to apply the settings immediately. This option is available for PERC 9 controllers only. If there are jobs to be completed, then this option is grayed-out. This job take at least 2 minutes to complete.
 - At Next Reboot Select this option to apply the settings during the next system reboot.
 - At Scheduled Time Select this option to apply the settings at a scheduled day and time:
 - **Start Time** and **End Time** Click the calendar icons and select the days. From the drop-down menus, select the time. The settings are applied between the start time and end time.
 - \circ $\,$ From the drop-down menu, select the type of reboot:
 - No Reboot (Manually Reboot System)
 - Graceful Shutdown
 - Force Shutdown
 - Power Cycle System (cold boot)
 - Add to Pending Operations Select this option to create a pending operation to apply the settings. You can view all pending operations for a controller in the Storage > Overview > Pending Operations page.
 - (i) NOTE:
 - The Add to Pending Operations option in not applicable for the Pending Operations page and for PCIe SSDs in the Physical Disks > Setup page.
 - Only the **Apply Now** option is available on the **Enclosure Setup** page.
- 3. Click Apply.

Based on the operation mode selected, the settings are applied.

Choosing operation mode using RACADM

To select the operation mode, use the jobqueue command.

For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Viewing and applying pending operations

You can view and commit all pending operations for the storage controller. All the settings are either applied at once, during the next reboot, or at a scheduled time based on the selected options. You can delete all the pending operations for a controller. You cannot delete individual pending operations.

Pending Operations are created on the selected components (controllers, enclosures, physical disks, and virtual disks).

Configuration jobs are created only on controller. In case of PCIe SSD, job is created on PCIe SSD disk and not on the PCIe Extender.

Viewing, applying, or deleting pending operations using web interface

- In the iDRAC web interface, go to Storage > Overview > Pending Operations. The Pending Operations page is displayed.
- 2. From the **Component** drop-down menu, select the controller for which you want to view, commit, or delete the pending operations.

The list of pending operations is displayed for the selected controller.



(i) NOTE:

- Pending operations are created for import foreign configuration, clear foreign configuration, security key operations, and encrypt virtual disks. But, they are not displayed in the **Pending Operations** page and in the Pending Operations pop-up message.
- Jobs for PCIe SSD cannot be created from the **Pending Operations** page
- **3.** To delete the pending operations for the selected controller, click **Delete All Pending Operations**.
- 4. From the drop-down menu, select one of the following and click Apply to commit the pending operations:
 - **Apply Now** Select this option to commit all the operations immediately. This option is available for PERC 9 controllers with the latest firmware versions.
 - At Next Reboot Select this option to commit all the operations during the next system reboot.
 - At Scheduled Time Select this option to commit the operations at a scheduled day and time.
 - Start Time and End Time Click the calendar icons and select the days. From the drop-down menus, select the time. The action is applied between the start time and end time.
 - From the drop-down menu, select the type of reboot:
 - No Reboot (Manually Reboot System)
 - Graceful Shutdown
 - Force Shutdown
 - Power Cycle System (cold boot)
- 5. If the commit job is not created, a message indicating that the job creation was not successful is displayed. Also, the message ID and the recommended response action are displayed.
- 6. If the commit job is created successfully, a message indicating that the job ID is created for the selected controller is displayed. Click **Job Queue** to view the progress of the job in the **Job Queue** page.

If the clear foreign configuration, import foreign configuration, security key operations, or encrypt virtual disk operations are in pending state, and if these are the only operations pending, then you cannot create a job from the **Pending Operations** page. You must perform any other storage configuration operation or use RACADM or WSMan to create the required configuration job on the required controller.

You cannot view or clear pending operations for PCle SSDs in the **Pending Operations** page. Use the racadm command to clear the pending operations for PCle SSDs.

Viewing and applying pending operations using RACADM

To apply pending operations, use the **jobqueue** command.

For more information, see the *iDRAC RACADM Command Line Reference Guide* available at **dell.com/idracmanuals**.

Storage devices — apply operation scenarios

Case 1: selected an apply operation (apply now, at next reboot, or at scheduled time) and there are no existing pending operations

If you have selected **Apply Now**, **At Next Reboot**, or **At Scheduled Time** and then clicked **Apply**, first the pending operation is created for the selected storage configuration operation.

- If the pending operation is successful and there are no prior existing pending operations, then the job is created. If the job is created successfully, a message indicating that the job ID is created for the selected device is displayed. Click **Job Queue** to view the progress of the job in the **Job Queue** page. If the job is not created, a message indicating that the job creation was not successful is displayed. Also, the message ID and the recommended response action are displayed.
- If the pending operation creation is unsuccessful and there are no prior existing pending operations, an error message with ID and recommended response action is displayed.

Case 2: selected an apply operation (apply now, at next reboot, or at scheduled time) and there are existing pending operations

If you have selected **Apply Now**, **At Next Reboot**, or **At Scheduled Time** and then clicked **Apply**, first the pending operation is created for the selected storage configuration operation.

• If the pending operation is created successfully and if there are existing pending operations, then a message is displayed.



- Click the View Pending Operations link to view the pending operations for the device.
- Click Create Job to create job for the selected device. If the job is created successfully, a message indicating that the job ID is created for the selected device is displayed. Click Job Queue to view the progress of the job in the Job Queue page. If the job is not created, a message indicating that the job creation was not successful is displayed. Also, the message ID and the recommended response action is displayed.
- Click Cancel to not create the job and remain on the page to perform more storage configuration operations.
- If the pending operation is not created successfully and if there are existing pending operations, then an error message is displayed.
- \circ $\;$ Click $\mbox{Pending Operations}$ to view the pending operations for the device.
- Click Create Job For Successful Operations to create the job for the existing pending operations. If the job is created successfully, a message indicating that the job ID is created for the selected device is displayed. Click Job Queue to view the progress of the job in the Job Queue page. If the job is not created, a message indicating that the job creation was not successful is displayed. Also, the message ID and the recommended response action are displayed.

• Click **Cancel** to not create the job and remain on the page to perform more storage configuration operations.

Case 3: selected add to pending operations and there are no existing pending operations

If you have selected **Add to Pending Operations** and then clicked **Apply**, first the pending operation is created for the selected storage configuration operation.

- If the pending operation is created successfully and if there are no existing pending operations, then an information message is displayed:
 - Click **OK** to remain on the page to perform more storage configuration operations.
 - Click **Pending Operations** to view the pending operations for the device. Until the job is created on the selected controller, these pending operations are not applied.
- If the pending operation is not created successfully and if there are no existing pending operations, then an error message is displayed.

Case 4: selected add to pending operations and there are prior existing pending operations

If you have selected **Add to Pending Operations** and then clicked **Apply**, first the pending operation is created for the selected storage configuration operation.

- If the pending operation is created successfully and if there are existing pending operations, then an information message is displayed:
 - $\circ~$ Click OK to remain on the page to perform more storage configuration operations.
 - Click **Pending Operations** to view the pending operations for the device.
- If the pending operation is not created successfully and if there are existing pending operations, then an error message is displayed.
 - Click **OK** to remain on the page to perform more storage configuration operations.
 - Click Pending Operations to view the pending operations for the device.

() NOTE:

- At any time, if you do not see the option to create a job on the storage configuration pages, go to Storage Overview > Pending Operations page to view the existing pending operations and to create the job on the required controller.
- Only cases 1 and 2 are applicable for PCIe SSD. You cannot view the pending operations for PCIe SSDs and hence Add
- to Pending Operations option is not available. Use racadm command to clear the pending operations for PCIe SSDs.

Blinking or unblinking component LEDs

You can locate a physical disk, virtual disk drive and PCIe SSDs within an enclosure by blinking one of the Light Emitting Diodes (LEDs) on the disk.

You must have Login privilege to blink or unblink an LED.

The controller must be real-time configuration capable. The real-time support of this feature is available only in PERC 9.1 firmware and later.

(i) NOTE: Blink or unblink is not supported for servers without backplane.

Blinking or unblinking component LEDs using web interface

To blink or unblink a component LED:



- 1. In the iDRAC Web interface, go to any of the following pages as per your requirement:
 - **Storage** > **Overview** > **Physical Disks** > **Status** Displays the identified Physical Disks page where you can blink or unblink the physical disks and PCIe SSDs.
 - Storage > Overview > Virtual Disks > Status- Displays the identified Virtual Disks page where you can blink or unblink the virtual disks.
- 2. If you select the physical disk:
 - Select or deselect all component LEDs Select the **Select/Deselect All** option and click **Blink** to start blinking the component LEDs. Similarly, click **Unblink** to stop blinking the component LEDs.
 - Select or deselect individual component LEDs Select one or more component(s) and click **Blink** to start blinking the selected component LED(s). Similarly, click **Unblink** to stop blinking the component LEDs.
- 3. If you select the virtual disk:
 - Select or deselect all physical disk drives or PCIe SSDs Select the **Select/Deselect All** option and click **Blink** to start blinking all the physical disk drives and the PCIe SSDs. Similarly, click **Unblink** to stop blinking the LEDs.
 - Select or deselect individual physical disk drives or PCIe SSDs Select one or more physical disk drives and click **Blink** to start blinking the LEDs for the physical disk drives or the PCIe SSDs. Similarly, click **Unblink** to stop blinking the LEDs.
- 4. If you are on the Identify Virtual Disk page:
 - Select or deselect all virtual disks Select the **Select/Deselect All** option and click **Blink** to start blinking the LEDs for all the virtual disks. Similarly, click **Unblink** to stop blinking the LEDs.
 - Select or deselect individual virtual disks Select one or more virtual disks and click **Blink** to start blinking the LEDs for the virtual disks. Similarly, click **Unblink** to stop blinking the LEDs.

If the blink or unblink operation is not successful, error messages are displayed.

Blinking or unblinking component LEDs using RACADM

To blink or unblink component LEDs, use the following commands:

racadm storage blink:<PD FQDD, VD FQDD, or PCIe SSD FQDD>

racadm storage unblink:<PD FQDD, VD FQDD, or PCIe SSD FQDD>

For more information, see the iDRAC RACADM Command Line Reference Guide available at dell.com/idracmanuals.

Warm reboot

When warm reboot is performed, following behaviors are observed:

- PERC controllers in iDRAC UI are grayed out immediately after warm reboot. They are available once re-inventory is completed after warm reboot. This is only applicable for PERC controllers and not for NVME/HBA/BOSS.
- Storage files in SupportAssist are empty when PERC controllers are grayed out in GUI.
- LC Logging for PAST event and Critical events are done for PERC during perc reinventory. Rest all LCL for PERC components are suppressed. LCL resumes after PERC re-inventory finishes.
- You cannot start any Real-time job until PERC re-inventory is finished.
- Telemetry data is not collected until PERC re-inventory is finished.
- After the PERC inventory is finished, the behavior is normal.



BIOS Settings

You can view multiple attributes, which are being used for a specific server under the BIOS Settings. You can modify different parameters of each attribute from this BIOS configuration setting. Once you select one attribute, it shows different parameters which are related to that specific attribute. You can modify multiple parameters of an attribute and apply changes before modifying a different attribute. When a user expands a configuration group, attributes are displayed in an alphabetical order.

() NOTE:

- Attribute level help content are dynamically generated.
- The iDRAC Direct USB port is available without host reboot, even when all USB ports are disabled.

Apply

Apply button remains greyed-out until any of the attributes are modified. Once you made changes to an attribute and click **Apply**, it allows you to modify the attribute with required changes. In case, the request fails to set the BIOS attribute, it throws an error with corresponding HTTP Response Status code mapped to SMIL API error or Job Creation error. A message is generated and displayed at that point. For more information, see *Event and Error Message Reference Guide for 14th Generation Dell EMC PowerEdge Servers* available at https://www.dell.com/idracmanuals.

Discard changes

The **Discard Changes** button is greyed-out until any of the attributes are modified. If you click **Discard Changes** button, all the recent changes are discarded and restored with the previous or initial values.

Apply and Reboot

When a user modifies value of an attribute or boot sequence, user is presented with two choices to apply the configuration; **Apply and Reboot** or **Apply on Next Reboot**. In either of the apply options, user is redirected to the job queue page to monitor the progress of that specific job.

A user can view auditing information related to BIOS configuration in the LC logs.

If you click **Apply and Reboot**, it restarts the server immediately to configure all the required changes. In case, the request fails to set the BIOS attributes, it throws an error with corresponding HTTP Response Status code mapped to SMIL API error or Job Creation error. An EEMI message is generated and displayed at that point.

Apply At Next Reboot

When a user modifies value of an attribute or boot sequence, user is presented with two choices to apply the configuration; **Apply and Reboot** or **Apply on Next Reboot**. In either of the apply options, user is redirected to the job queue page to monitor the progress of that specific job.

A user can view auditing information related to BIOS configuration in the LC logs.

If you click **Apply At Next Reboot**, it configures all the required changes on the next restart of the server. You will not experience any immediate modifications based on the recent configuration changes until the next reboot session is taking place successfully. In case, the request fails to set the BIOS attributes, it throws an error with corresponding HTTP Response Status code mapped to SMIL API error or Job Creation error. An EEMI message is generated and displayed at that point.



Delete All Pending Values

Delete All pending Values button is enabled only when there are pending values based on the recent configuration changes. In case, user decides not to apply the configuration changes, user can click **Delete All Pending Values** button to terminate all the modifications. In case, the request fails to remove the BIOS attributes, it throws an error with corresponding HTTP Response Status code mapped to SMIL API error or Job Creation error. An EEMI message is generated and displayed at that point.

Pending Value

Configuration of a BIOS attribute via iDRAC is not applied immediately to BIOS. It requires a server reboot for the changes to take place. When you modify a BIOS attribute then **Pending Value** gets updated. If an attribute already has a pending value (and that has been configured) it is displayed on the GUI.

Modifying BIOS Configuration

Modifying BIOS configuration results in audit log entries, which gets entered in LC logs.

BIOS Live Scanning

BIOS live scanning verifies the integrity and authenticity of the BIOS image in the BIOS primary ROM when the host is powered ON but not in POST.

() NOTE:

- This feature requires iDRAC Datacenter license.
- You need to have Debug privilege for operating this feature.

iDRAC performs verification of immutable sections of BIOS image automatically at the followings scenarios:

- At AC cycle/Cold boot
- On a schedule determined by user
- On demand (initiated by user)

Successful result of live scanning is logged to LC log. Failure result is logged to both LCL and SEL.

Topics:

- BIOS Live Scanning
- BIOS Recovery and Hardware Root of Trust (RoT)

BIOS Live Scanning

BIOS live scanning verifies the integrity and authenticity of the BIOS image in the BIOS primary ROM when the host is powered ON but not in POST.

() NOTE:

- This feature requires iDRAC Datacenter license.
- You need to have Debug privilege for operating this feature.

iDRAC performs verification of immutable sections of BIOS image automatically at the followings scenarios:

- At AC cycle/Cold boot
- On a schedule determined by user
- On demand (initiated by user)

Successful result of live scanning is logged to LC log. Failure result is logged to both LCL and SEL.



BIOS Recovery and Hardware Root of Trust (RoT)

For PowerEdge server, it is mandatory to recover from corrupted or damaged BIOS image either due to malicious attack or power surges or any other unforeseeable events. An alternate reserve of BIOS image would be necessary to recover BIOS in order to bring the PowerEdge server back to functional mode from unbootable mode. This alternative/recovery BIOS is stored in a 2nd SPI (mux'ed with primary BIOS SPI).

The recovery sequence can be initiated through any of the following approaches with iDRAC as the main orchestrator of the BIOS recovery task:

- 1. Auto recovery of BIOS primary image/recovery image BIOS image is recovered automatically during the host boot process after the BIOS corruption is detected by BIOS itself.
- 2. Forced recovery of BIOS Primary/recovery image User initiates an OOB request to update BIOS either because they have a new updated BIOS or BIOS was just crashing by failing to boot.
- **3. Primary BIOS ROM update** The single Primary ROM is split into Data ROM and Code ROM. iDRAC has full access/ control over Code ROM. It switches MUX to access Code ROM whenever needed.
- 4. BIOS Hardware Root of Trust (RoT) This feature is available in severs with model number RX5X, CX5XX, and TX5X. During every host boot (only cold boot or A/C cycle, not during warm reboot), iDRAC ensures that RoT is performed. RoT runs automatically and user cannot initiate it using any interfaces. This iDRAC boot first policy verifies host BIOS ROM contents on every AC cycle and host DC cycle. This process ensures secure boot of BIOS and further secures the host boot process.

NOTE: For more information on Hardware RoT, refer to this link: https://downloads.dell.com/Manuals/Common/dell-emcidrac9-security-root-of-trust-bios-live-scanning.pdf



Configuring and using virtual console

iDRAC has added an enhanced HTML5 option in vConsole which allows vKVM (virtual Keyboard, Video, and Mouse) over standard VNC client. You can use the virtual console to manage a remote system using the keyboard, video, and mouse on your management station to control the corresponding devices on a managed server. This is a licensed feature for rack and tower servers. It is available by default in blade servers. You need iDRAC Configure privilege to access all configurations on virtual console.

Following are the list of configurable attributes in Virtual Console:

- vConsole Enabled Enabled / Disabled
- Max Sessions 1-6
- Active sessions 0-6
- Remote Presence Port (Not applicable for eHTML5 plugin)
- Video Encryption Enabled / Disabled (Not applicable for eHTML5 plugin)
- Local Server Video Enabled / Disabled
- Plug-in Type eHTML5 (by default), ActiveX, Java, HTML5
- Dynamic Action on Sharing Request Timeout Full Access, Read Only Access, And Deny Access
- Automatic System Lock Enabled / Disabled
- Keyboard/Mouse Attach State Auto-attach, Attached, and Detached

The key features are:

- A maximum of six simultaneous Virtual Console sessions are supported. All the sessions view the same managed server console simultaneously.
- You can launch virtual console in a supported web browser by using Java, ActiveX, HTML5, or eHTML5 plug-in.
 NOTE: By default, the virtual console type is set to eHTML5.

(i) NOTE: Any change in web server configuration will result in termination of existing virtual console session.

- When you open a Virtual Console session, the managed server does not indicate that the console has been redirected.
- You can open multiple Virtual Console sessions from a single management station to one or more managed systems simultaneously.
- You cannot open two virtual console sessions from the management station to the managed server using the same HTML5 plug-in.
- If a second user requests a Virtual Console session, the first user is notified and is given the option to refuse access, allow read-only access, or allow full shared access. The second user is notified that another user has control. The first user must respond within thirty seconds, or else access is granted to the second user based on the default setting. If neither the first or second user has administrator privileges, terminating the first user's session automatically terminates the second user's session.
- Boot logs and crash logs are captured as Video logs and are in MPEG1 format.
- Crash screen is captured as JPEG file.
- Keyboard macros are supported on all plug-ins.
- Keyboard macros are supported on all plug-ins. Following are the list of macros that are supported by ActiveX and Java plug-ins:

Table 57. Keyboard Macros Supported by ActiveX and Java plug-ins

MAC Client	Win Client	Linux Client
Ctrl-Al-Del	Ctrl-Alt-Del	Ctrl-Alt-Del
Alt-SysRq-B	Alt-SysRq-B	Alt-SysRq-B
-	Win-P	-
-	-	Ctrl-Alt-F<1-12>
Alt-SysRq	-	-



Table 57. Keyboard Macros Supported by ActiveX and Java plug-ins (continued)

MAC Client	Win Client	Linux Client
SysRq	-	-
PrtScrn	-	-
Alt-PrtScrn	-	-
Pause	-	-

(i) **NOTE:** For keyboard macros supported in HTML plug-in, see the section HTML5 based virtual console.

NOTE: The number of active virtual-console sessions displayed in the web interface is only for active web-interface sessions. This number does not include sessions from other interfaces such as SSH and RACADM.

NOTE: For information about configuring your browser to access the virtual console, see Configuring web browsers to use virtual console.

(i) NOTE: To disable KVM access, use the **Disable** option under the settings for chassis in the OME Modular web interface.

Topics:

- Supported screen resolutions and refresh rates
- Configuring virtual console
- Previewing virtual console
- Launching virtual console
- Using virtual console viewer

Supported screen resolutions and refresh rates

The following table lists the supported screen resolutions and corresponding refresh rates for a Virtual Console session running on the managed server.

Table 58. Supported screen resolutions and refresh rates

Screen Resolution	Refresh Rate (Hz)
720x400	70
640x480	60, 72, 75, 85
800×600	60, 70, 72, 75, 85
1024x768	60, 70, 72, 75, 85
1280x1024	60
1920x1200	60

It is recommended that you configure the monitor display resolution to 1920x1200 pixels.

Virtual Console supports a maximum video resolution of 1920x1200 at 60 Hz refresh rate. In order to achieve this resolution, following conditions are required:

- KVM / monitor attached to VGA that supports 1920x1200 resolution
- Latest Matrox video driver (for Windows)

When a local KVM / Monitor with maximum resolution below 1920x1200 is connected to either VGA connector, it will reduce the maximum resolution supported in virtual console.

iDRAC virtual console leverages the onboard Matrox G200 graphics controller to determine the maximum resolution of the attached monitor when a physical display is present. When the monitor supports 1920x1200 or greater resolution, the virtual



console supports 1920x1200 resolution. If the monitor attached supports lower max resolution (like many KVMs), the virtual console max resolution is limited.

Maximum virtual console resolutions based on monitor display ratio:

- 16:10 monitor: 1920x1200 will be the max resolution
- 16:9 monitor: 1920x1080 will be the max resolution

When a physical monitor is not connected to either VGA port on the server, the OS installed will dictate the available resolutions for virtual console.

Maximum virtual console resolutions based on host OS without physical monitor:

- Windows: 1600x1200 (1600x1200, 1280x1024, 1152x864, 1024x768, 800x600)
- Linux: 1024x768 (1024x768, 800x600, 848x480, 640x480)

NOTE: If a higher resolution through virtual console is required when physical KVM or monitor is not present, a VGA Display Emulator dongle can be leveraged to mimic an external monitor connection with a resolution up to 1920x1080.

(i) NOTE: If you have an active Virtual Console session and a lower resolution monitor is connected to the Virtual Console, the server console resolution may reset if the server is selected on the local console. If the system is running a Linux operating system, an X11 console may not be viewable on the local monitor. Press <Ctrl><Alt><F1> at the iDRAC Virtual Console to switch Linux to a text console.

Configuring virtual console

Before configuring the Virtual Console, make sure that the management station is configured.

You can configure the virtual console using iDRAC Web interface or RACADM command line interface.

Configuring virtual console using web interface

To configure Virtual Console using iDRAC Web interface:

- 1. Go to Configuration > Virtual Console. Click Start the Virtual Console link, then Virtual Console page is displayed.
- 2. Enable virtual console and specify the required values. For information about the options, see the *iDRAC Online Help*.
 (i) NOTE: If you are using Nano operating system, disable the Automatic System Lock feature on the Virtual Console
- 3. Click Apply. The virtual console is configured.

page.

Configuring virtual console using RACADM

To configure the Virtual Console, use the set command with the objects in the **iDRAC.VirtualConsole** group.

For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Previewing virtual console

Before launching the Virtual Console, you can preview the state of the Virtual Console on the **System** > **Properties** > **System Summary** page. The **Virtual Console Preview** section displays an image showing the state of the Virtual Console. The image is refreshed every 30 seconds. This is a licensed feature.

(i) NOTE: The Virtual Console image is available only if you have enabled Virtual Console.

Launching virtual console

You can launch the virtual console using the iDRAC Web Interface or a URL.

(i) NOTE: Do not launch a Virtual Console session from a Web browser on the managed system.



Before launching the Virtual Console, make sure that:

- You have administrator privileges.
- Web browser is configured to use HTML5, eHTML5, Java, or ActiveX plug-ins.
- Minimum network bandwidth of 1 MB/sec is available.

NOTE: If the embedded video controller is disabled in BIOS and if you launch the Virtual Console, the Virtual Console Viewer is blank.

While launching Virtual Console using 32-bit or 64-bit IE browsers, use HTML5/eHTML5, or use the required plug-in (Java or ActiveX) that is available in the respective browser. The Internet Options settings are common for all browsers.

While launching the Virtual Console using Java plug-in, occasionally you may see a Java compilation error. To resolve this, go to **Java control panel** > **General** > **Network Settings** and select **Direct Connection**.

If the Virtual Console is configured to use ActiveX plug-in, it may not launch the first time. This is because of the slow network connection and the temporary credentials (that Virtual Console uses to connect) timeout is two minutes. The ActiveX client plug-in download time may exceed this time. After the plug-in is successfully downloaded, you can launch the Virtual Console normally.

To launch the Virtual Console by using HTML5/eHTML5 plug-in, you must disable the pop-up blocker.

Virtual Console has the following Console controls:

- 1. General You can set Keyboard Macros, Aspect Ratio, and Touch Mode.
- 2. ${\rm KVM}$ Shows the values for Frame Rate, Bandwidth, Compression, and Packet Rate.
- **3. Performance** You can change the Video quality and Video speed using this option.
- 4. User List You can view the list of users connected to the console.

You can access Virtual Media by clicking the Connect to Virtual Media option available in virtual console.

Launching virtual console using web interface

You can launch the virtual console in the following ways:

• Go to Configuration > Virtual Console. Click Start the Virtual Console link. Virtual console page is displayed.

The **Virtual Console Viewer** displays the remote system's desktop. Using this viewer, you can control the remote system's mouse and keyboard functions from your management station.

Multiple message boxes may appear after you launch the application. To prevent unauthorized access to the application, navigate through these message boxes within three minutes. Otherwise, you are prompted to relaunch the application.

If one or more Security Alert windows appear while launching the viewer, click Yes to continue.

Two mouse pointers may appear in the viewer window: one for the managed server and another for your management station.

Launching virtual console using a URL

To launch the Virtual Console using the URL:

- 1. Open a supported Web browser and in the address box, type the following URL in lower case: https://iDRAC_ip/console
- 2. Based on the login configuration, the corresponding **Login** page is displayed:
 - If Single Sign On is disabled and Local, Active Directory, LDAP, or Smart Card login is enabled, the corresponding **Login** page is displayed.
 - If Single-Sign On is enabled, the Virtual Console Viewer is launched and the Virtual Console page is displayed in the background.
 - () NOTE: Internet Explorer supports Local, Active Directory, LDAP, Smart Card (SC) and Single Sign-On (SSO) logins. Firefox supports Local, AD, and SSO logins on Windows-based operating system and Local, Active Directory, and LDAP logins on Linux-based operating systems.
 - **NOTE:** If you do not have Access Virtual Console privilege but have Access Virtual Media privilege, then using this URL launches the Virtual Media instead of the Virtual Console.



Disabling warning messages while launching virtual console or virtual media using Java or ActiveX plug-in

You can disable the warning messages while launching the Virtual Console or Virtual Media using Java plug-in.

(i) NOTE: You need Java 8 or later to use this feature and to launch iDRAC Virtual Console over an IPv6 network.

1. Initially, when you launch Virtual Console or Virtual Media using Java plug-in, the prompt to verify the publisher is displayed. Click **Yes**.

A certificate warning message is displayed indicating that a trusted certificate is not found.

NOTE: If the certificate is found in the operating system's certificate store or if it is found in a previously specified user location, then this warning message is not displayed.

2. Click Continue.

The Virtual Console Viewer or Virtual Media Viewer is launched.

(i) NOTE: The Virtual Media viewer is launched if Virtual Console is disabled.

- 3. From the Tools menu, click Session Options and then Certificate tab.
- 4. Click Browse Path, specify the location to store the user's certificate, click Apply, click OK, and exit from the viewer.
- 5. Launch Virtual Console again.
- 6. In the certificate warning message, select the Always trust this certificate option, and then click Continue.
- 7. Exit from the viewer.
- 8. When you re-launch Virtual Console, the warning message is not displayed.

Using virtual console viewer

The Virtual Console Viewer provides various controls such as mouse synchronization, virtual console scaling, chat options, keyboard macros, power actions, next boot devices, and access to Virtual Media. For information to use these features, see the *iDRAC Online Help*.

(i) NOTE: If the remote server is powered off, the message 'No Signal' is displayed.

The Virtual Console Viewer title bar displays the DNS name or the IP address of the iDRAC you are connected to from the management station. If iDRAC does not have a DNS name, then the IP address is displayed. The format is:

• For rack and tower servers:

<DNS name / IPv6 address / IPv4 address>, <Model>, User: <username>, <fps>

• For blade servers:

<DNS name / IPv6 address / IPv4 address>, <Model>, <Slot number>, User: <username>, <fps>

Sometimes the Virtual Console Viewer may display low quality video. This is due to slow network connectivity that leads to loss of one or two video frames when you start the Virtual Console session. To transmit all the video frames and improve the subsequent video quality, do any of the following:

- In the System Summary page, under Virtual Console Preview section, click Refresh.
- In the Virtual Console Viewer, under Performance tab, set the slider to Maximum Video Quality.

eHTML5 based virtual console

NOTE: While using eHTML5 to access virtual console, the language must be consistent across client and target keyboard layout, OS, and browser. For example, all must be in English (US) or any of the supported languages.

To launch the eHTML5 virtual console, you must enable the virtual console feature from the iDRAC Virtual Console page and set the **Plug-in Type** option to eHTML5.

(i) NOTE: By default the virtual console type is set to eHTML5.

You can launch virtual console as a pop-up window by using one of the following methods:

• From iDRAC Home page, click the Start the Virtual Console link available in the Console Preview session

278 Configuring and using virtual console



- From iDRAC Virtual Console page, click Start the Virtual Console link.
- From iDRAC login page, type https//<iDRAC IP>/console. This method is called as Direct Launch.

In the eHTML5 virtual console, the following menu options are available:

- Power
- Boot
- Chat
- Keyboard
- Screen Capture
- Refresh
- Full Screen
- Disconnect Viewer
- Console Controls
- Virtual Media

The **Pass all keystrokes to server** option is not supported on eHTML5 virtual console. Use keyboard and keyboard macros for all the functional keys.

- General
 - **Console control** This has the following configuration options:
 - Keyboard Macros This is supported in eHTML5 virtual console and are listed as the following drop-down options. Click **Apply** to apply the selected key combination on the server.
 - Ctrl+Alt+Del
 - Ctrl+Alt+F1
 - Ctrl+Alt+F2
 - Ctrl+Alt+F3
 - Ctrl+Alt+F4
 - Ctrl+Alt+F5
 - Ctrl+Alt+F6
 - Ctrl+Alt+F7
 - Ctrl+Alt+F8
 - Ctrl+Alt+F9
 - Ctrl+Alt+F10
 - Ctrl+Alt+F11
 - Ctrl+Alt+F12
 - Alt+Tab
 - Alt+ESC
 - Ctrl+ESC
 - Alt+Space
 - Alt+Enter
 - Alt+Hyphen
 - Alt+F1
 - Alt+F2
 - Alt+F3
 - Alt+F4
 - Alt+F5
 - Alt+F6
 - Alt+F7
 - Alt+F8
 - Alt+F9
 - Alt+F10
 - Alt+F11
 - Alt+F12
 - PrntScrn
 - Alt+PrntScrn
 - F1
 - Pause
 - Tab
 - Ctrl+Enter



- SysRq
- Alt+SysRq
- Win-P
- Aspect Ratio The eHTML5 virtual console video image automatically adjusts the size to make the image visible. The following configuration options are displayed as a drop-down list:
 - Maintain
 - Don't Maintain

Click **Apply** to apply the selected settings on the server.

- Touch Mode The eHTML5 virtual console supports the Touch Mode feature. The following configuration options are displayed as a drop-down list:
 - Direct
 - Relative

Click **Apply** to apply the selected settings on the server.

- **Virtual Clipboard** Virtual clipboard enables you to cut / copy / paste text buffer from virtual console to iDRAC host server. Host server could be BIOS, UEFI or in OS prompt. This is a one-way action from client computer to iDRAC's host server only. Follow these steps to use the Virtual clipboard:
 - Place the mouse cursor or keyboard focus on the desired window in the host server desktop.
 - Select the **Console Controls** menu from vConsole.
 - Copy the OS clipboard buffer using keyboard hotkeys, mouse, or touch pad controls depending on the Client OS. Or, you can type the text manually in the text box.

• Click Send Clipboard to Host.

 \circ $\;$ Then, the text appears on the host server's active window.

(i) NOTE:

- This feature is only available in Datacenter license.
- This feature only supports ASCII text.
- Control characters are not supported.
- Characters such as **New line** and **Tab** are allowed.
- Text buffer size is limited to 4000 characters.
- If more than maximum buffer is pasted, then the edit box in iDRAC GUI will truncate it to maximum buffer size.
- **KVM** This menu has list of the following read only components:
- Frame Rate
- Bandwidth
- Compression
- Packet Rate
- Performance You can use the slider button to adjust Maximum Video Quality and Maximum Video Speed.
- User List You can see the list of users that are logged in to the Virtual console.
- **Keyboard** The difference between physical and virtual keyboard is that virtual keyboard changes its layout according to the browser language.
- Virtual Media Click Connect Virtual Media option to start the virtual media session.
- **Connect Virtual Media** This menu contains the options for Map CD/DVD, Map Removable Disk, Map External Device, and Reset USB.
 - **Virtual Media Statistics** This menu shows the Transfer Rate (Read-only). Also, it shows the details of CD/DVD and Removable Disks details such as Mapping details, status (read-only or not), duration, and Read/Write Bytes.
- **Create Image** This menu allows you to select a local folder and generate FolderName.img file with local folder contents.

NOTE: For security reasons read/write access is disabled while accessing virtual console in eHTML5. With Java or ActiveX plug-ins, you can accept security messaging before the plug-in is given the read/write authority.

Supported Browsers

The eHTML5 virtual console is supported on the following browsers:

- Internet Explorer 11
- Chrome 78/79
- Firefox 70/71


• Safari 13.1

(i) NOTE: It is recommended to have Mac OS version 10.10.2 (or onward) installed in the system.

For more details on supported browsers and versions, see the *iDRAC Release Notes* available at https://www.dell.com/ idracmanuals.

HTML5 based virtual console

() **NOTE:** While using HTML5 to access virtual console, the language must be consistent across client and target keyboard layout, OS, and browser. For example, all must be in English (US) or any of the supported languages.

To launch the HTML5 virtual console, you must enable the virtual console feature from the iDRAC Virtual Console page and set the **Plug-in Type** option to HTML5.

You can launch virtual console as a pop-up window by using one of the following methods:

- From iDRAC Home page, click the Start the Virtual Console link available in the Console Preview session
- From iDRAC Virtual Console page, click Start the Virtual Console link.
- From iDRAC login page, type https//<iDRAC IP>/console. This method is called as Direct Launch.

In the HTML5 virtual console, the following menu options are available:

- Power
- Boot
- Chat
- Keyboard
- Screen Capture
- Refresh
- Full Screen
- Disconnect Viewer
- Console Controls
- Virtual Media

The **Pass all keystrokes to server** option is not supported on HTML5 virtual console. Use keyboard and keyboard macros for all the functional keys.

- **Console control** This has the following configuration options:
 - Keyboard Macros This is supported in HTML5 virtual console and are listed as the following drop-down options. Click **Apply** to apply the selected key combination on the server.
 - Ctrl+Alt+Del
 - Ctrl+Alt+F1
 - Ctrl+Alt+F2
 - Ctrl+Alt+F3
 - Ctrl+Alt+F4
 - Ctrl+Alt+F5
 - Ctrl+Alt+F6
 - Ctrl+Alt+F7
 - Ctrl+Alt+F8
 - Ctrl+Alt+F9
 - Ctrl+Alt+F10
 - Ctrl+Alt+F11
 - Ctrl+Alt+F12
 - Alt+Tab
 - Alt+ESC
 - Ctrl+ESC
 - Alt+Space
 - Alt+Enter
 - Alt+Hyphen
 - Alt+F1
 - Alt+F2
 - Alt+F3

Configuring and using virtual console 281



- Alt+F4
- Alt+F5
- Alt+F6
- Alt+F7
- Alt+F8
- Alt+F9
- Alt+F10
- Alt+F11
- Alt+F12
- PrntScrn
- Alt+PrntScrn
- F1
- Pause
- Tab
- Ctrl+Enter
- SysRq
- Alt+SysRq
- Win-P
- Aspect Ratio The HTML5 virtual console video image automatically adjusts the size to make the image visible. The following configuration options are displayed as a drop-down list:
 - Maintain
 - Don't Maintain

Click **Apply** to apply the selected settings on the server.

- Touch Mode The HTML5 virtual console supports the Touch Mode feature. The following configuration options are displayed as a drop-down list:
 - Direct
 - Relative

Click **Apply** to apply the selected settings on the server.

- **Virtual Clipboard** Virtual clipboard enables you to cut / copy / paste text buffer from virtual console to iDRAC host server. Host server could be BIOS, UEFI or in OS prompt. This is a one-way action from client computer to iDRAC's host server only. Follow these steps to use the Virtual clipboard:
 - Place the mouse cursor or keyboard focus on the desired window in the host server desktop.
 - Select the **Console Controls** menu from vConsole.
 - Copy the OS clipboard buffer using keyboard hotkeys, mouse, or touch pad controls depending on the Client OS. Or, you can type the text manually in the text box.
 - Click Send Clipboard to Host.
 - \circ $\;$ Then, the text appears on the host server's active window.

(i) NOTE:

- This feature is only available in Datacenter license.
- This feature only supports ASCII text.
- Control characters are not supported.
- Characters such as **New line** and **Tab** are allowed.
- Text buffer size is limited to 4000 characters.
- If more than maximum buffer is pasted, then the edit box in iDRAC GUI will truncate it to maximum buffer size.
- **Keyboard** The difference between physical and virtual keyboard is that virtual keyboard changes its layout according to the browser language.
- Touch Mode The HTML5 virtual console supports the Touch Mode feature. The following configuration options are displayed as a drop-down list:
 - o Direct

Relative

- Click **Apply** to apply the selected settings on the server.
- Mouse Acceleration Select the mouse acceleration based on the operating system. The following configuration options are displayed as a drop-down list:
 - Absolute (Windows, latest versions of Linux, Mac OS-X)



- Relative, no acceleration
- Relative (RHEL, earlier versions of Linux)
- Linux RHEL 6.x and SUSE Linux Enterprise Server 11 or later

Click **Apply** to apply the selected settings on the server.

• Virtual Media — Click Connect Virtual Media option to start the virtual media session. when the virtual media is connected, you can see the options like Map CD/DVD, Map Removable Disk, and Reset USB.

NOTE: For security reasons read/write access is disabled while accessing virtual console in HTML5. With Java or ActiveX plug-ins, you can accept security messaging before the plug-in is given the read/write authority.

Supported Browsers

The HTML5 virtual console is supported on the following browsers:

- Internet Explorer 11
- Chrome 78/79
- Firefox 70/71
- Safari 13.1

(i) NOTE: It is recommended to have Mac OS version 10.10.2 (or onward) installed in the system.

For more details on supported browsers and versions, see the *iDRAC Release Notes* available at https://www.dell.com/ idracmanuals.

Synchronizing mouse pointers

(i) **NOTE:** This feature is not applicable with eHTML5 plugin type.

When you connect to a managed system through the Virtual Console, the mouse acceleration speed on the managed system may not synchronize with the mouse pointer on the management station and displays two mouse pointers in the Viewer window.

When using Red Hat Enterprise Linux or Novell SUSE Linux, configure the mouse mode for Linux before you launch the Virtual Console viewer. The operating system's default mouse settings are used to control the mouse arrow in the Virtual Console viewer.

When two mouse cursors are seen on the client Virtual Console viewer, it indicates that the server's operating system supports Relative Positioning. This is typical for Linux operating systems or Lifecycle Controller and causes two mouse cursors if the server's mouse acceleration settings are different from the mouse acceleration settings on the Virtual Console client. To resolve this, switch to single cursor or match the mouse acceleration on the managed system and the management station:

- To switch to single cursor, from the **Tools** menu, select **Single Cursor**.
- To set the mouse acceleration, go to Tools > Session Options > Mouse. Under Mouse Acceleration tab, select Windows
 or Linux based on the operating system.

To exit single cursor mode, press <F9> or the configured termination key.

NOTE: This is not applicable for managed systems running Windows operating system since they support Absolute Positioning.

When using the Virtual Console to connect to a managed system with a recent Linux distribution operating system installed, you may experience mouse synchronization problems. This may be due to the Predictable Pointer Acceleration feature of the GNOME desktop. For correct mouse synchronization in the iDRAC Virtual Console, this feature must be disabled. To disable Predictable Pointer Acceleration, in the mouse section of the **/etc/X11/xorg.conf** file, add:

Option "AccelerationScheme" "lightweight".

If synchronization problems continue, do the following additional change in the <user_home>/.gconf/desktop/gnome/ peripherals/mouse/%gconf.xml file:

Change the values for motion_threshold and motion_acceleration to -1.

If you turn off mouse acceleration in GNOME desktop, in the Virtual Console viewer, go to **Tools** > **Session Options** > **Mouse**. Under **Mouse Acceleration** tab, select **None**.

For exclusive access to the managed server console, you must disable the local console and re-configure the **Max Sessions** to 1 on the **Virtual Console page**.

Configuring and using virtual console 283



Passing all keystrokes through virtual console for Java or ActiveX plug-in

You can enable the **Pass all keystrokes to server** option and send all keystrokes and key combinations from the management station to the managed system through the Virtual Console Viewer. If it is disabled, it directs all the key combinations to the management station where the Virtual Console session is running. To pass all keystrokes to the server, in the Virtual Console Viewer, go to **Tools** > **Session Options** > **General** tab and select the **Pass all keystrokes to server** option to pass the management station's keystrokes to the managed system.

The behavior of the Pass all keystrokes to server feature depends on the:

Plug-in type (Java or ActiveX) based on which Virtual Console session is launched.

For the Java client, the native library must be loaded for Pass all keystrokes to server and Single Cursor mode to function. If the native libraries are not loaded, the **Pass all keystrokes to server** and **Single Cursor** options are deselected. If you attempt to select either of these options, an error message is displayed indicating that the selected options are not supported.

For the ActiveX client, the native library must be loaded for Pass all keystrokes to server function to work. If the native libraries are not loaded, the **Pass all keystrokes to server** option is deselected. If you attempt to select this option, an error message is displayed indicating that the feature is not supported

For MAC operating systems, enable the **Enable access of assistive device** option in **Universal Access** for the Pass all keystrokes to server feature to work.

- Operating system running on the management station and managed system. The key combinations that are meaningful to the operating system on the management station are not passed to the managed system.
- Virtual Console Viewer mode—Windowed or Full Screen.

In Full Screen mode, **Pass all keystrokes to server** is enabled by default.

In Windowed mode, the keys passed only when the Virtual Console Viewer is visible and is active.

When changed from Full Screen mode to Windowed mode, the previous state of Pass all keys is resumed.

Java-based virtual console session running on Windows operating system

- Ctrl+Alt+Del key is not sent to the managed system, but always interpreted by the management station.
 - When Pass All Keystrokes to Server is enabled, the following keys are not sent to the managed system:
 - Browser Back Key
 - Browser Forward Key
 - Browser Refresh key
 - o Browser Stop Key
 - Browser Search Key
 - Browser Favorites key
 - Browser Start and Home key
 - Volume mute key
 - Volume down key
 - Volume up key
 - Next track key
 - Previous track key
 - Stop Media key
 - Play/Pause media key
 - Start mail key
 - Select media key
 - Start Application 1 key
 - Start Application 2 key
- All the individual keys (not a combination of different keys, but a single key stroke) are always sent to the managed system. This includes all the Function keys, Shift, Alt, Ctrl key and Menu keys. Some of these keys affect both management station and managed system.

For example, if the management station and the managed system is running Windows operating system, and Pass All Keys is disabled, when you press the Windows key to open the **Start** Menu, the **Start** menu opens on both management station and



managed system. However, if Pass All Keys is enabled, then the **Start** menu is opened only on the managed system and not on the management station.

• When Pass All Keys is disabled, the behavior depends on the key combinations pressed and the special combinations interpreted by the operating system on the management station.

Java based virtual console session running on Linux operating system

The behavior mentioned for Windows operating system is also applicable for Linux operating system with the following exceptions:

- When Pass all keystrokes to server is enabled, <Ctrl+Alt+Del> is passed to the operating system on the managed system.
- Magic SysRq keys are key combinations interpreted by the Linux Kernel. It is useful if the operating system on the
 management station or the managed system freezes and you need to recover the system. You can enable the magic SysRq
 keys on the Linux operating system using one of the following methods:
 - Add an entry to **/etc/sysctl.conf**
 - echo "1" > /proc/sys/kernel/sysrq
- When Pass all keystrokes to server is enabled, the magic SysRq keys are sent to the operating system on the managed system. The key sequence behavior to reset the operating system, that is reboot without un-mounting or sync, depends on whether the magic SysRq is enabled or disabled on the management station:
 - If SysRq is enabled on the management station, then <Ctrl+Alt+SysRq+b> or <Alt+SysRq+b> resets the management station irrespective of the system's state.
 - If SysRq is disabled on the management station, then the <Ctrl+Alt+SysRq+b> or <Alt+SysRq+b>keys resets the operating system on the managed system.
 - Other SysRq key combinations (example, <Alt+SysRq+k>, <Ctrl+Alt+SysRq+m>, and so on) are passed to the managed system irrespective of the SysRq keys enabled or not on the management station.

Using SysRq magic keys through remote console

You can enable SysRq magic keys through the remote console using any of the following:

- Opensoure IPMI tool
- Using SSH or External Serial Connector

Using opensource IPMI tool

Make sure that BIOS/iDRAC settings supports console redirection using SOL.

1. At the command prompt, run the SOL activate command:

Ipmitool -I lanplus -H <ipaddr> -U <username> -P <passwd> sol activate

The SOL session is activated.

- 2. After the server boots to the operating system, the localhost.localdomain login prompt appears. Log in using the operating system user name and password.
- **3.** If SysRq is not enabled, enable using echo 1 >/proc/sys/kernel/sysrq.
- 4. Run break sequence ~B.
- **5.** Use the SysRq magic key to enable the SysRq function. For example, the following command displays the memory information on the console:

echo m > /proc/sysrq-trigger displays

Using SSH or external serial connector directly connecting through serial cable

- 1. For SSH sessions, after logging in using the iDRAC username and password, at the /admin> prompt, run the command console com2. The localhost.localdomain prompt appears.
- 2. For console redirection using external serial connector directly connected to the system through a serial cable, the localhost.localdomain login prompt appears after the server boots to the operating system.
- 3. Log in using the operating system user name and password.
- 4. If SysRq is not enabled, enable using echo 1 >/proc/sys/kernel/sysrq.



5. Use the magic key to enable the SysRg function. For example, the following command reboots the server:

echo b > /proc/sysrq-trigger

(i) NOTE: You do not have to run break sequence before using the magic SysRq keys.

ActiveX based virtual console session running on Windows operating system

The behavior of the pass all keystrokes to server feature in ActiveX based Virtual Console session running on Windows operating system is similar to the behavior explained for Java based Virtual Console session running on the Windows management station with the following exceptions:

• When Pass All Keys is disabled, pressing F1 launches the application Help on both management station and managed system, and the following message is displayed:

Click Help on the Virtual Console page to view the online Help

- The media keys may not be blocked explicitly.
- <Alt + Space>, <Ctrl + Alt + +>, <Ctrl + Alt + -> are not sent to the managed system and is interpreted by the operating system on the management station.



Using iDRAC Service Module

The iDRAC Service Module is a software application that is recommended to be installed on the server (it is not installed by default). It complements iDRAC with monitoring information from the operating system. It complements iDRAC by providing additional data to work with iDRAC interfaces such as the Web interface, Redfish, RACADM, and WSMan. You can configure the features monitored by the iDRAC Service Module to control the CPU and memory consumed on the server's operating system. Host OS command line interface has been introduced to enable or disable status of Full Power Cycle for all System components except the PSU.

(i) NOTE: iDRAC9 uses iSM version 3.01 and higher.

NOTE: You can use the iDRAC Service Module only if you have installed iDRAC Express or iDRAC Enterprise/Datacenter license.

Before using iDRAC Service Module, ensure that:

- You have login, configure, and server control privileges in iDRAC to enable or disable the iDRAC Service Module features.
- You do not disable the **iDRAC Configuration using local RACADM** option.
- OS to iDRAC pass-through channel is enabled through the internal USB bus in iDRAC.

(i) NOTE: If you perform LC wipe, idrac.Servicemodule values may still show the old values.

() NOTE:

- When iDRAC Service Module runs for the first time, by default it enables the OS to iDRAC pass-through channel in iDRAC. If you disable this feature after installing the iDRAC Service Module, then you must enable it manually in iDRAC.
- If the OS to iDRAC pass-through channel is enabled through LOM in iDRAC, then you cannot use the iDRAC Service Module.

Topics:

- Installing iDRAC Service Module
- Supported operating systems for iDRAC Service Module
- iDRAC Service Module monitoring features
- Using iDRAC Service Module from iDRAC web interface
- Using iDRAC Service Module from RACADM

Installing iDRAC Service Module

You can download and install the iDRAC Service Module from **dell.com/support**. You must have administrator privilege on the server's operating system to install the iDRAC Service Module. For information on installation, see the iDRAC Service Module User's Guide available at www.dell.com/idracservicemodule.

(i) NOTE: This feature is not applicable for Dell Precision PR7910 systems.

Installing iDRAC Service Module from iDRAC Express and Basic

From the iDRAC Service Module Setup page, click Install Service Module.

- 1. The Service Module Installer is available to the host operating system and a job is created in iDRAC. For Microsoft Windows operating system or Linux operating system, log in to the server either remotely or locally.
- 2. Find the mounted volume labeled as "SMINST" on your device list and run the appropriate script:
 - On Windows, open the command prompt and run the **ISM-Win.bat** batch file.
 - On Linux, open the shell prompt and run the ISM-Lx.sh script file.
- 3. After the installation is complete, iDRAC displays the Service Module as Installed and the installation date.



NOTE: The installer will be available to the host operating system for 30 minutes. If you do not start the installation within 30 minutes, you must restart the Service Module installation.

Installing iDRAC Service Module from iDRAC Enterprise

- 1. On the SupportAssist Registration wizard, click Next.
- 2. On the iDRAC Service Module Setup page, click Install Service Module.
- 3. Click Launch Virtual Console and click Continue on the security warning dialog box.
- 4. To locate the iSM installer file, log in to the server either remotely or locally.

(i) **NOTE:** The installer will be available to the host operating system for 30 minutes. If you do not start the installation within 30 minutes, you must restart the installation.

- 5. Find the mounted volume labeled as "SMINST" on your device list and run the appropriate script:
 - On Windows, open the command prompt and run the **ISM-Win.bat** batch file.
 - On Linux, open the shell prompt and run the **ISM-Lx.sh** script file.
- Follow the instructions on the screen to complete the installation. On the iDRAC Service Module Setup page, the Install Service Module button is disabled after the installation is complete and the Service Module status is displayed as Running.

Supported operating systems for iDRAC Service Module

For the list of operating systems supported by the iDRAC Service Module, see the iDRAC Service Module User's Guide available at www.dell.com/idracservicemodule.

iDRAC Service Module monitoring features

The iDRAC Service Module (iSM) provides the following monitoring features:

- Redfish profile support for network attributes
- iDRAC Hard Reset
- iDRAC access via Host OS (Experimental Feature)
- In-band iDRAC SNMP alerts
- View operating system (OS) information
- Replicate Lifecycle Controller logs to operating system logs
- Perform automatic system recovery options
- Populate Windows Management Instrumentation (WMI) Management Providers
- Integrate with SupportAssist Collection. This is applicable only if iDRAC Service Module version 2.0 or later is installed.
- Prepare to Remove NVMe PCIe SSD. For more informationhttps://www.dell.com/support/article/sln310557 .
- Remote Server Power Cycle

Redfish profile support for network attributes

iDRAC Service Module v2.3 or later provides additional network attributes to iDRAC, which can be obtained through the REST clients from iDRAC. For more details, see iDRAC Redfish profile support.

Operating system information

The OpenManage Server Administrator currently shares operating system information and host name with iDRAC. The iDRAC Service Module provides similar information such as OS name, OS version, and Fully Qualified Domain Name (FQDN) with



iDRAC. By default, this monitoring feature is enabled. It is not disabled if OpenManage Server Administrator is installed on the host OS.

In iSM version 2.0 or later, the operating system information feature is amended with the OS network interface monitoring. When iDRAC Service Module version 2.0 or later is used with iDRAC 2.00.00.00, it starts monitoring the operating system network interfaces. You can view this information using iDRAC web interface, RACADM, or WSMan.

Replicate Lifecycle logs to OS log

You can replicate the Lifecycle Controller Logs to the OS logs from the time when the feature is enabled in iDRAC. This is similar to the System Event Log (SEL) replication performed by OpenManage Server Administrator. All events that have the **OS Log** option selected as the target (in the **Alerts** page, or in the equivalent RACADM or WSMan interfaces) are replicated in the OS log using the iDRAC Service Module. The default set of logs to be included in the OS logs is the same as configured for SNMP alerts or traps.

iDRAC Service Module also logs the events that have occurred when the operating system is not functioning. The OS logging performed by iDRAC Service Module follows the IETF syslog standards for Linux-based operating systems.

() NOTE: Starting iDRAC Service Module version 2.1, the Lifecycle Controller Logs replication location in the Windows OS logs can be configured using the iDRAC Service Module installer. You can configure the location while installing iDRAC Service Module or modifying the iDRAC Service Module installer.

If OpenManage Server Administrator is installed, this monitoring feature is disabled to avoid duplicate SEL entries in the OS log.

NOTE: On Microsoft Windows, if iSM events get logged under System logs instead of Application logs, restart the Windows Event Log service or restart the host OS.

Automatic system recovery options

The Automatic system recovery feature is a hardware-based timer. If a hardware failure occurs, a notification may not be available, but the server is reset as if the power switch was activated. ASR is implemented using a timer that continuously counts down. The Health Monitor frequently reloads the counter to prevent it from counting down to zero. If the ASR counts down to zero, it is assumed that the operating system has locked up and the system automatically attempts to reboot.

You can perform automatic system recovery operations such as reboot, power cycle, or power off the server after a specified time interval. This feature is enabled only if the operating system watchdog timer is disabled. If OpenManage Server Administrator is installed, this monitoring feature is disabled to avoid duplicate watchdog timers.

Windows Management Instrumentation providers

WMI is a set of extensions to the Windows Driver Model that provides an operating system interface through which instrumented components provide information and notification. WMI is Microsoft's implementation of the Web-Based Enterprise Management (WBEM) and Common Information Model (CIM) standards from the Distributed Management Task Force (DMTF) to manage Server hardware, operating systems and applications. WMI Providers helps to integrate with Systems Management Consoles such as Microsoft System Center and enables scripting to manage Microsoft Windows Servers.

You can enable or disable the WMI option in iDRAC. iDRAC exposes the WMI classes through the iDRAC Service Module providing the server's health information. By default, WMI information feature is enabled. The iDRAC Service Module exposes the WSMan monitored classes in iDRAC through WMI. The classes are exposed in the root/cimv2/dcim namespace.

The classes can be accessed using any of the standard WMI client interfaces. For more information, see the profile documents.

This content use the **DCIM_iDRACCardString** and **DCIM_iDRACCardInteger** classes to illustrate the capability that WMI information feature provides in iDRAC Service Module. For the details of the supported classes and profiles, see the WSMan profiles documentation available at https://www.dell.com/support.

The listed attributes are used to configure **User Accounts** along with the required privileges:

AttributeName	WSMAN-Class	Privilege	License	Description	Supported Operation
UserName	DCIM_iDRACCardS tring	Write Privileges: ConfigUsers, Login	Basic	16users:	Enum, Get, Invoke

AttributeName	WSMAN-Class	Privilege	License	Description	Supported Operation
		Read Privileges:Login		Users.1#UserName to Users.16#UserNam e	
Password	DCIM_iDRACCardS tring	Write Privileges: ConfigUsers, Login Read Privileges:	Basic	Users.1#Password to Users.16#Password	Enum, Get, Invoke
Privilege	DCIM_iDRACCardI nteger	Vrite Privileges:	Basic Users.1#Password to Users.16#Passwor	Users.1#Password	Enum, Get, Invoke
		ConfigUsers, Login Read Privileges: Login		Users.16#Password	

- Enumerate or Get operation on the mentioned classes will provide the attribute related data.
- The attribute can be set by invoking the ApplyAttribute or SetAttribute command from **DCIM_iDRACCardService** class.
- () NOTE: DCIM_Account class is removed from WSMAN and provided the feature through attribute model.
 - **DCIM_iDRACCardString** and **DCIM_iDRACCardInteger** classes provide similar support to configure iDRAC users accounts.

Remote iDRAC Hard Reset

By using iDRAC, you can monitor the supported servers for critical system hardware, firmware, or software issues. Sometimes, iDRAC may become unresponsive due to various reasons. During such scenarios, you must turn off the server and reset iDRAC. To reset the iDRAC CPU, you must either power off and power on the server or perform an AC power cycle.

By using the remote iDRAC hard reset feature, whenever iDRAC becomes unresponsive, you can perform a remote iDRAC reset operation without an AC power cycle.. To reset the iDRAC remotely, make sure that you have administrative privileges on the host OS. By default, the remote iDRAC hard reset feature is enabled. You can perform a remote iDRAC hard reset using iDRAC Web interface, RACADM, and WSMan.

Command usage

This section provides the command usages for Windows, Linux, and ESXi operating systems to perform iDRAC hard reset.

Windows

- Using the local Windows Management Instrumentation (WMI):
- o winrm i iDRACHardReset wmi/root/cimv2/dcim/DCIM_iSMService? InstanceID="iSMExportedFunctions"
- Using the remote WMI interface:

```
winrm i iDRACHardReset wmi/root/cimv2/dcim/dcim_ismservice?
InstanceID="iSMExportedFunctions" -u:<admin-username> -p:<admin-password> -r:http://
<remote-hostname OR IP>/wsman -a:Basic -encoding:utf-8 -skipCACheck -skipCNCheck
```

• Using the Windows PowerShell script with force and without force:

```
Invoke-iDRACHardReset -force
```

Invoke-iDRACHardReset

• Using the **Program Menu** shortcut:

For simplicity, iSM provides a shortcut in the **Program Menu** of the Windows operating system. When you select the **Remote iDRAC Hard Reset** option, you are prompted for a confirmation to reset the iDRAC. After you confirm, the iDRAC is reset and the result of the operation is displayed.

NOTE: The following warning message appears in the **Event Viewer** under the **Application Logs** category. This warning does not require any further action.



() NOTE: A provider, ismserviceprovider, has been registered in the Windows Management Instrumentation namespace Root\CIMV2\DCIM to use the LocalSystem account. This account is privileged and the provider may cause a security violation if it does not correctly impersonate user requests.

• Linux

iSM provides an executable command on all iSM supported Linux operating system. You can run this command by logging into the operating system by using SSH or equivalent.

Invoke-iDRACHardReset

Invoke-iDRACHardReset -f

ESXi

On all iSM supported ESXi operating systems, the iSM v2.3 supports a Common Management Programming Interface (CMPI) method provider to perform the iDRAC reset remotely by using the WinRM remote commands.

```
winrm i iDRACHardReset http://schemas.dell.com/wbem/wscim/1/cim-schema/2/root/cimv2/dcim/
DCIM_iSMService?__cimnamespace=root/cimv2/dcim+InstanceID= iSMExportedFunctions -u:<root-
username> -p:<passwd> -r:https://<Host-IP>:443/WSMan -a:basic -encoding:utf-8
-skipCNCheck -skipCACheck -skipRevocationcheck
```

(i) NOTE: VMware ESXi operating system does not prompt for confirmation before resetting the iDRAC.

NOTE: Due to limitations on the VMware ESXi operating system, iDRAC connectivity is not restored completely after the reset. Ensure that you manually reset iDRAC.

Table 59. Error Handling

Result	Description	
0	Success	
1	Unsupported BIOS version for iDRAC reset	
2	Unsupported platform	
3	Access denied	
4	iDRAC reset failed	

In-band Support for iDRAC SNMP Alerts

By using iDRAC Service Module v2.3, you can receive SNMP alerts from the host operating system, which is similar to the alerts that are generated by iDRAC.

You can also monitor the iDRAC SNMP alerts without configuring the iDRAC and manage the server remotely by configuring the SNMP traps and destination on the host OS. In iDRAC Service Module v2.3 or later, this feature converts all the Lifecycle logs replicated in the OS logs into SNMP traps.

(i) NOTE: This feature is active only if the Lifecycle Logs replication feature is enabled.

NOTE: On Linux operating systems, this feature requires a master or OS SNMP enabled with SNMP multiplexing (SMUX) protocol.

By default, this feature is disabled. Though the In-band SNMP alerting mechanism can coexist along with iDRAC SNMP alerting mechanism, the recorded logs may have redundant SNMP alerts from both the sources. It is recommended to either use the in-band or out-of-band option, instead of using both.

Command usage

This section provides the command usages for Windows, Linux, and ESXi operating systems.

• Windows operating system

• Using the local Windows Management Instrumentation (WMI):

```
winrm i EnableInBandSNMPTraps wmi/root/cimv2/dcim/DCIM_iSMService?
InstanceID="iSMExportedFunctions" @{state="[0/1]"}
```



• Using the remote WMI interface:

```
winrm i EnableInBandSNMPTraps wmi/root/cimv2/dcim/DCIM_iSMService?
InstanceID="iSMExportedFunctions" @{state="[0/1]"} -u:<admin-username> -p:<admin-
passwd> -r:http://<remote-hostname OR IP>/WSMan -a:Basic -encoding:utf-8 -skipCACheck
-skipCNCheck
```

Linux operating system

On all iSM supported Linux operating system, iSM provides an executable command. You can run this command by logging into the operating system by using SSH or equivalent.

Beginning with iSM 2.4.0, you can configure Agent-x as the default protocol for in-band iDRAC SNMP alerts using the following command:

./Enable-iDRACSNMPTrap.sh 1/agentx -force

- If -force is not specified, ensure that the net-SNMP is configured and restart the snmpd service.
- To enable this feature:

```
Enable-iDRACSNMPTrap.sh 1
```

```
Enable-iDRACSNMPTrap.sh enable
```

• To disable this feature:

```
Enable-iDRACSNMPTrap.sh 0
```

Enable-iDRACSNMPTrap.sh disable

NOTE: The **--force** option configures the Net-SNMP to forward the traps. However, you must configure the trap destination.

VMware ESXi operating system

On all iSM supported ESXi operating systems, the iSM v2.3 supports a Common Management Programming Interface (CMPI) method provider to enable this feature remotely by using the WinRM remote commands.

```
winrm i EnableInBandSNMPTraps http://schemas.dell.com/wbem/
wscim/1/cim-schema/2/root/cimv2/dcim/DCIM_iSMService? __cimnamespace=root/cimv2/
dcim+InstanceID=iSMExportedFunctions -u:<user-name> -p:<passwd> -r:https://<remote-host-
name</pre>
```

ip-address>:443/WSMan -a:basic -encoding:utf-8 -skipCNCheck -skipCACheck -skipRevocationcheck @{state="[0/1]"}

(i) NOTE: You must review and configure the VMware ESXi system-wide SNMP settings for traps.

NOTE: For more details, refer to the **In-BandSNMPAlerts** technical white paper available at https://www.dell.com/ support.

iDRAC access via Host OS

By using this feature, you can configure and monitor the hardware parameters through iDRAC Web interface, WSMan, and RedFish interfaces using the host IP address without configuring the iDRAC IP address. You can use the default iDRAC credentials if the iDRAC server is not configured or continue to use the same iDRAC credentials if the iDRAC server was configured earlier.

iDRAC access via Windows Operating Systems

You can perform this task by using the following methods:

- Install the iDRAC access feature by using the webpack.
- Configure using iSM PowerShell script

Installation by using MSI

292 Using iDRAC Service Module



You can install this feature by using the web-pack. This feature is disabled on a typical iSM installation. If enabled, the default listening port number is 1266. You can modify this port number within the range 1024 through 65535. iSM redirects the connection to the iDRAC. iSM then creates an inbound firewall rule, OS2iDRAC. The listening port number is added to the OS2iDRAC firewall rule in the host operating system, which allows incoming connections. The firewall rule is enabled automatically when this feature is enabled.

Beginning with iSM 2.4.0, you can retrieve the current status and listening-port configuration by using the following Powershell cmdlet:

Enable-iDRACAccessHostRoute -status get

The output of this command indicates whether this feature is enabled or disabled. If the feature is enabled, it displays the listening-port number.

(i) NOTE: Ensure that the Microsoft IP Helper Services is running on your system for this feature to function.

To access the iDRAC Web interface, use the format https://<host-name>or OS-IP>:443/login.html in the browser, where:

- <host-name> Complete host name of the server on which iSM is installed and configured for iDRAC access via OS feature. You can use the OS IP address if the host name is not present.
- 443 Default iDRAC port number. This is called the Connect Port number to which all the incoming connections on listen port number are redirected. You can modify the port number through iDRAC Web interface, WSMan, and RACADM interfaces.

Configuration by using iSM PowerShell cmdlet

If this feature is disabled while installing iSM, you can enable the feature by using the following Windows PowerShell command provided by iSM:

Enable-iDRACAccessHostRoute

If the feature is already configured, you can disable or modify it by using the PowerShell command and the corresponding options. The available options are as follows:

- Status This parameter is mandatory. The values are not case sensitive and the value can be true, false, or get.
- **Port** This is the listening port number. If you do not provide a port number, the default port number (1266) is used. If the **Status** parameter value is FALSE, then you can ignore rest of the parameters. You must enter a new port number that is not already configured for this feature. The new port number settings overwrite the existing OS2iDRAC in-bound firewall rule and you can use the new port number to connect to iDRAC. The value range is from 1024 to 65535.
- **IPRange** This parameter is optional and it provides a range of IP addresses that are allowed to connect to iDRAC through the host operating system. The IP address range format is in Classless Inter-Domain Routing (CIDR) format, which is a combination of IP address and subnet mask. For example, 10.94.111.21/24. Access to iDRAC is restricted for IP addresses that are not within the range.

(i) NOTE: This feature supports only IPv4 addresses.

iDRAC access via Linux Operating Systems

You can install this feature by using the setup.sh file that is available with the web-pack. This feature is disabled on a default or typical iSM installation. To get the status of this feature, use the following command:

Enable-iDRACAccessHostRoute get-status

To install, enable, and configure this feature, use the following command:

./Enable-iDRACAccessHostRoute <Enable-Flag> [<source-port> <source-IP-range/source-iprange-mask>]

<Enable-Flag>=0

Disable

<source-port> and <source-IP-range/source-ip-range-mask> are not required.

<Enable-Flag>=1

Enable

<source-port> is required and <source-ip-range-mask> is optional.

<source-IP-range>

Using iDRAC Service Module 293



Coexistence of OpenManage Server Administrator and iDRAC Service Module

In a system, both OpenManage Server Administrator and the iDRAC Service Module can co-exist and continue to function correctly and independently.

If you have enabled the monitoring features during the iDRAC Service Module installation, then after the installation is complete if the iDRAC Service Module detects the presence of OpenManage Server Administrator, it disables the set of monitoring features that overlap. If OpenManage Server Administrator is running, the iDRAC Service Module disables the overlapping monitoring features after logging to the OS and iDRAC.

When you re-enable these monitoring features through the iDRAC interfaces later, the same checks are performed and the features are enabled depending on whether OpenManage Server Administrator is running or not.

Using iDRAC Service Module from iDRAC web interface

To use the iDRAC Service Module from the iDRAC web interface:

- Go to IDRAC Settings > Overview > iDRAC Service Module > Configure Service Module. The iDRAC Service Module Setup page is displayed.
- 2. You can view the following:
 - Installed iDRAC Service Module version on the host operating system
 - Connection status of the iDRAC Service Module with iDRAC.
 - () NOTE: When a server has multiple operating systems and iDRAC Service Module is installed in all operating systems, then iDRAC connects only with the most recent instance of iSM among all operating systems. An error is displayed for all the older instances of iSM on other operating systems. To connect iSM with iDRAC on any other operating system which already has iSM installed, uninstall and reinstall iSM on that particular operating system.
- 3. To perform out-of-band monitoring functions, select one or more of the following options:
 - **OS Information**—View the operating system information.
 - **Replicate Lifecycle Log in OS Log**—Include Lifecycle Controller logs to operating system logs. This option is disabled if OpenManage Server Administrator is installed on the system.
 - **WMI Information** Include WMI information.
 - Auto System Recovery Action—Perform auto recovery operations on the system after a specified time (in seconds):
 - Reboot
 - Power Off System
 - Power Cycle System

This option is disabled if OpenManage Server Administrator is installed on the system.

Using iDRAC Service Module from RACADM

To use the iDRAC Service Module from RACADM, use the objects in the ServiceModule group.

For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.



Using USB port for server management

On the 14th generation servers, a dedicated micro USB port is available to configure iDRAC. You can perform the following functions using the micro USB port:

- Connect to the system using the USB network interface to access system management tools such as iDRAC web interface and RACADM.
- Configure a server by using SCP files that are stored on a USB drive.
- () NOTE: To manage a USB port or to configure a server by importing Server Configuration Profile (SCP) files on a USB drive, you must have the System Control privilege.

(i) NOTE: An alert / report is generated when a USB device is inserted. This feature is only available on Intel based servers.

To configure Management USB Settings, go to **iDRAC Settings** > **Settings** > **Management USB Settings**. Following options are available:

• **USB Management Port**—Select **Enabled** to enable the port either to import the SCP file when a USB drive is connected or to access iDRAC using the micro USB port.

(i) NOTE: Ensure that the USB drive contains a valid SCP file.

(i) NOTE: Use an OTG adapter to convert from Type-A to Micro-B USB. Connections from USB hubs are not supported.

- **iDRAC Managed: USB SCP**—Select from following options to configure the system by importing SCP stored on a USB drive:
 - **Disabled**—Disables SCP imports
 - **Enabled only when server has default credential settings** If this option is selected then the SCP can only be imported when the default password is not changed for the following:
 - BIOS
 - iDRAC web interface
 - Enabled only for compressed configuration files—Select this option to allow SCP file import only if the files are in compressed format.

NOTE: Selecting this option allows you to password protect the compressed file. You can enter a password to secure the file by using **Password for Zip file** option.

• Enabled—Select this option to allow importing SCP file without running a check during runtime.

Topics:

- Accessing iDRAC interface over direct USB connection
- Configuring iDRAC using server configuration profile on USB device

Accessing iDRAC interface over direct USB connection

The iDRAC direct feature allows you to directly connect your laptop to the iDRAC USB port. This feature allows you to interact directly with the iDRAC interfaces such as the web interface, RACADM, and WSMan for advanced server management and servicing.

For a list of supported browsers and operating systems, see the *iDRAC Release Notes* available at https://www.dell.com/ idracmanuals.

(i) NOTE: If you are using Windows operating system, you may need to install an RNDIS driver to use this feature.

To access the iDRAC interface over the USB port:

- 1. Turn off any wireless networks and disconnect from any other hard wired network.
- 2. Ensure that the USB port is enabled. For more information, see Configuring USB management port settings.

Using USB port for server management 295



- **3.** Wait for the laptop to acquire IP address 169.254.0.4. It may take several seconds for the IP addresses to be acquired. iDRAC acquires the IP address 169.254.0.3.
- 4. Start using iDRAC network interfaces such as the web interface, RACADM, Redfish or WSMan. For example, to access the iDRAC web interface, open a supported browser, and type the address 169.254.0.3 and press enter.
- 5. When iDRAC is using the USB port, the LED blinks indicating activity. The blink frequency is four per second.
- **6.** After completing the desired actions, disconnect the USB cable from the system. The LED turns off.

Configuring iDRAC using server configuration profile on USB device

With the iDRAC USB management port, you can configure iDRAC at-the-server. Configure the USB Management port settings in iDRAC, insert the USB device that has the server configuration profile, and then import the server configuration profile from the USB device to iDRAC.

NOTE: You can set the USB Management port settings using the iDRAC interfaces only if there is no USB device connected to the server.

Configuring USB management port settings

You can enable or disable the iDRAC Direct USB Port using the system BIOS. Navigate to **System BIOS** > **Integrated Devices**. Select **On** to enable and **Off** to disable the iDRAC Direct USB Port.

In iDRAC, you must have Server Control privilege to configure the USB management port. When a USB device is connected, the **System Inventory** page displays the USB device information under the Hardware Inventory section.

An event is logged in the Lifecycle Controller logs when:

- The device is in Automatic or iDRAC mode, and USB device is inserted or removed.
- USB Management Port Mode is modified.
- Device is automatically switched from iDRAC to OS.
- Device is ejected from iDRAC or OS

When a device exceeds its power requirements as allowed by USB specification, the device is detached and an over-current event is generated with the following properties:

- Category : System Health
- Type: USB device
- Severity: Warning
- Allowed notifications: Email, SNMP trap, remote syslog, and WS-Eventing
- Actions: None

An error message is displayed and logged to Lifecycle Controller log when:

- You try to configure the USB management port without the Server Control user privilege.
- A USB device is in use by iDRAC, and you attempt to modify the USB Management Port Mode.
- A USB device is in use by iDRAC, and you remove the device.

Configuring USB management port using web interface

To configure the USB port:

- 1. In the iDRAC Web interface, go to **iDRAC Settings** > **Settings** > **Management USB Settings**.
- 2. The USB Management Port is set to Enabled.
- **3.** From the **iDRAC Managed: USB SCP** Configuration drop-down menu, select options to configure a server by importing Server Configuration Profile files stored on a USB drive:
 - Disabled
 - Enabled only when server has default credential settings
 - Enabled only for compressed configuration files
 - Enabled



For information about the fields, see the *iDRAC Online Help*.

() NOTE: iDRAC9 allows you to password protect the compressed file after you select Enabled only for compressed configuration files to compress the file before importing. You can enter a password to secure the file by using Password for Zip file option.

4. Click **Apply** to apply the settings.

Configuring USB management port using RACADM

To configure the USB management port, use the following RACADM sub commands and objects:

• To view the USB port status:

racadm get iDRAC.USB.PortStatus

• To view the USB port configuration:

racadm get iDRAC.USB.ManagementPortMode

• To view USB device inventory:

racadm hwinventory

• To set up overcurrent alert configuration:

```
racadm eventfilters
```

For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Configuring USB management port using iDRAC settings utility

To configure the USB port:

- In the iDRAC Settings Utility, go to Media and USB Port Settings. The iDRAC Settings Media and USB Port Settings page is displayed.
- 2. From the **iDRAC Direct: USB Configuration XML** drop-down menu, select options to configure a server by importing server configuration profile stored on a USB drive:
 - Disabled
 - Enabled while server has default credential settings only
 - Enabled only for compressed configuration files
 - Enabled
 - For information about the fields, see the *iDRAC Settings Utility Online Help*.
- 3. Click **Back**, click **Finish** and then click **Yes** to apply the settings.

Importing Server Configuration Profile from USB device

Make sure to create a directory in root of the USB device called System_Configuration_XML that contains both the config and control files:

- Server Configuration Profile (SCP) is in the System_Configuration_XML sub-directory under the USB device root directory. This file includes all the attribute-value pairs of the server. This includes attributes of iDRAC, PERC, RAID, and BIOS. You can edit this file to configure any attribute on the server. The file name can be <servicetag>-config.xml, <servicetag>-config.json, <modelnumber>-config.xml, <modelnumber>-config.json, config.yml or config.json.
- Control file Includes parameters to control the import operation and does not have attributes of iDRAC or any other component in the system. The control file contain three parameters:
 - ShutdownType Graceful, Forced, No Reboot.
 - TimeToWait (in secs) 300 minimum and 3600 maximum.
 - EndHostPowerState on/off.

Example of control.xml file:



<InstructionTable> <InstructionRow> <InstructionType>Configuration XML import Host control Instruction </InstructionType> <Instruction>ShutdownType</Instruction> <Value>NoReboot</Value> <ValuePossibilities>Graceful,Forced,NoReboot</ValuePossibilities> </InstructionRow> <InstructionRow> <InstructionType>Configuration XML import Host control Instruction </InstructionType> <Instruction>TimeToWait</Instruction> <Value>300</Value> <ValuePossibilities>Minimum value is 300 -Maximum value is 3600 seconds.</ValuePossibilities> </InstructionRow> <InstructionRow> <InstructionType>Configuration XML import Host control Instruction </InstructionType> <Instruction>EndHostPowerState</Instruction> <Value>On</Value> <ValuePossibilities>On,Off</ValuePossibilities> </InstructionRow> </InstructionTable>

You must have Server Control privilege to perform this operation.

NOTE: While importing the SCP, changing the USB management settings in the SCP file results in a failed job or job completed with errors. You can comment out the attributes in the SCP to avoid the errors.

To import the server configuration profile from the USB device to iDRAC:

1. Configure the USB management port:

- Set USB Management Port Mode to Automatic or iDRAC.
- Set iDRAC Managed: USB XML Configuration to Enabled with default credentials or Enabled.
- 2. Insert the USB key (that has the configuration.xml and the control.xml file) to the iDRAC USB port.

(i) NOTE: File name and file type are case sensitive for XML files. Ensure that both are in lower case.

- **3.** The server configuration profile is discovered on the USB device in the System_Configuration_XML sub-directory under the USB device root directory. It is discovered in the following sequence:
 - <servicetag>-config.xml/<servicetag>-config.json
 - <modelnum>-config.xml/<modelnum>-config.json
 - config.xml/config.json
- 4. A server configuration profile import job starts.

If the profile is not discovered, then the operation stops.

If **iDRAC Managed: USB XML Configuration** was set to **Enabled with default credentials** and the BIOS setup password is not null or if one of the iDRAC user accounts have been modified, an error message is displayed and the operation stops.

- 5. LCD panel and LED, if present, display the status that an import job has started.
- 6. If there is a configuration that needs to be staged and the **Shut Down Type** is specified as **No Reboot** is specified in the control file, you must reboot the server for the settings to be configured. Else, server is rebooted and the configuration is applied. Only when the server was already powered down, then the staged configuration is applied even if the **No Reboot** option is specified.
- 7. After the import job is complete, the LCD/LED indicates that the job is complete. If a reboot is required, LCD displays the job status as "Paused waiting on reboot".
- 8. If the USB device is left inserted on the server, the result of the import operation is recorded in the results.xml file in the USB device.



LCD messages

If the LCD panel is available, it displays the following messages in a sequence:

- 1. Importing When the server configuration profile is being copied from the USB device.
- 2. Applying When the job is in-progress.
- **3.** Completed When the job has completed successfully.
- 4. Completed with errors When the job has completed with errors.
- 5. Failed When the job has failed.

For more details, see the results file on the USB device.

LED blinking behavior

The USB LED indicates the status of a server-configuration profile operation being performed using the USB port. The LED may not be available on all systems.

- Solid green The server configuration profile is being copied from the USB device.
- Blinking green The job is in progress.
- Blinking amber The job has failed or completed with errors.
- Solid green The job has completed successfully.

NOTE: On PowerEdge R840 and R940xa, if there is a LCD present, the USB LED does not blink when an import operation is in progress using the USB port. Check the status of the operation using the LCD.

Logs and results file

The following information is logged for the import operation:

- Automatic import from USB is logged in the Lifecycle Controller log file.
- If the USB device is left inserted, the job results are recorded in the Results file located in the USB key.
- A Result file named Results.xml is updated or created in the subdirectory with the following information:
- Service tag Data is recorded after the import operation has either returned a job ID or returned an error.
- Job ID Data is recorded after the import operation has returned a job ID.
- Start Date and Time of Job Data is recorded after the import operation has returned a job ID.
- Status Data is recorded when the import operation returns an error or when the job results are available.



Using Quick Sync 2

With Dell OpenManage Mobile running on an Android or iOS mobile device, you can easily access server directly or through OpenManage Essentials or OpenManage Enterprise (OME) console. It allows you to review server details and inventory, view LC and System Event logs, get automatic notifications on mobile device from an OME console, assign IP address and modify iDRAC password, configure key BIOS attributes, and take remediation actions as needed. You can also power cycle a server, access system console, or access iDRAC GUI.

OMM can be downloaded for free from the Apple App Store, or from Google Play Store.

You must install the OpenManage Mobile application on the mobile device (supports Android 5.0+ and iOS 9.0+ mobile devices) to manage server using iDRAC Quick Sync 2 interface.

(i) NOTE: This section is displayed only in those servers that has Quick Sync 2 module in left rack ear.

(i) **NOTE:** This feature is currently supported on mobile devices with Android operating system and Apple iOS.

In the current release, this feature is available on all 14th generation of PowerEdge servers. It requires Quick Sync 2 Left Control Panel (embedded in **Left rack ear**) and Bluetooth Low Energy (and optionally Wi-Fi) enabled mobile devices. Therefore, it is a hardware up-sell and the feature capabilities are not dependent on iDRAC software licensing.

NOTE: For more information on configuring Quick Sync 2 in MX platform systems, see the OpenManage Enterprise Modular User's Guide and OpenManage Mobile User's Guide available at **dell.com/support/manuals**.

The iDRAC Quick Sync 2 Configuration procedures:

(i) NOTE: Not applicable for MX platforms.

Once Quick Sync is configured, activate the Quick Sync 2 button on the Left Control Panel. Make sure the Quick Sync 2 light turns on. Access the Quick Sync 2 Information using a mobile device (Android 5.0+ or IOS 9.0+, OMM 2.0 or above).

Using OpenManage Mobile, you can:

- View inventory information
- View monitoring information
- Configure the basic iDRAC network settings

For more information about OpenManage Mobile, see the *Dell EMC OpenManage Mobile User's Guide* available at https://www.dell.com/openmanagemanuals.

Topics:

- Configuring iDRAC Quick Sync 2
- Using mobile device to view iDRAC information

Configuring iDRAC Quick Sync 2

Using iDRAC web interface, RACADM, WSMan and iDRAC HII you can configure iDRAC Quick Sync 2 feature to allow access to the mobile device:

- Access Configure to Read-Write, Read-only, and Disabled. Read-Write is the default option.
- Timeout Configure to Enabled or Disabled. Enabled is the default option.
- **Timeout Limit** Indicates the time after which the Quick Sync 2 mode is disabled. By default, seconds are selected. The default value is 120 seconds. The range is 120 to 3600 seconds.
 - 1. If enabled, you can specify a time after which the Quick Sync 2 mode is turned off. To turn on, press the activation button again.
 - 2. If disabled, the timer does not allow you to enter a time-out period.
- **Read Authentication** Configures to Enabled, this is the default option.
- WiFi Configures to Enabled, this is the default option.



You must have Server Control privilege to configure the settings. A server reboot is not required for the settings to take effect. once configured, you can activate the Quick Sync 2 button on the Left Control Panel. Make sure the Quick Sync light turns on. Then, access the Quick Sync Information via a mobile device.

An entry is logged to the Lifecycle Controller log when the configuration is modified.

Configuring iDRAC Quick Sync 2 settings using web interface

To configure iDRAC Quick Sync 2:

- 1. In the iDRAC web interface, go to Configuration > System Settings > Hardware Settings > iDRAC Quick Sync.
- 2. In the **iDRAC Quick Sync** section, from the **Access** menu, select one of the following to provide access to the Android or iOS mobile device:
 - Read-write
 - Read-only
 - Disabled
- 3. Enable the Timer.
- **4.** Specify the Timeout Limit.
 - For more information about the fields, see the *iDRAC Online Help*.
- 5. Click **Apply** to apply the settings.

Configuring iDRAC Quick Sync 2 settings using RACADM

To configure the iDRAC Quick Sync 2feature, use the racadm objects in the **System.QuickSync** group. For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Configuring iDRAC Quick Sync 2 settings using iDRAC settings utility

To configure iDRAC Quick Sync 2:

- 1. In the iDRAC GUI, go to Configuration > Systems Settings > Hardware Settings > iDRAC Quick Sync.
- 2. In the iDRAC Quick Sync section:
 - Specify the access level.
 - Enable Timeout.
 - Specify the User Defined Timeout Limit (the range is 120 to 3600 seconds.).

For more information about the fields, see the *iDRAC Online Help*.

3. Click **Back**, click **Finish**, and then click **Yes**. The settings are applied.

Using mobile device to view iDRAC information

To view iDRAC information from the mobile device, see the *Dell EMC OpenManage Mobile User's Guide* available at https://www.dell.com/openmanagemanuals for the steps.



Managing virtual media

iDRAC provides virtual media with HTML5 based client with local ISO and IMG file, remote ISO and IMG file support. Virtual media allows the managed server to access media devices on the management station or ISO CD/DVD images on a network share as if they were devices on the managed server. You need iDRAC Configure privilege to modify the configuration.

Following are the configurable attributes:

- Attached Media Enabled Enabled / Disabled
- Attach Mode Auto-attach, Attached, and Detached
- Max Sessions 1
- Active Sessions 1
- Virtual Media Encryption Enabled (by default)
- Floppy Emulation Disabled (by default)
- Boot Once Enabled / Disabled
- Connection Status Connected / Disconnected

Using the Virtual Media feature, you can:

- Remotely access media connected to a remote system over the network
- Install applications
- Update drivers
- Install an operating system on the managed system

This is a licensed feature for rack and tower servers. It is available by default for blade servers.

The key features are:

- Virtual Media supports virtual optical drives (CD/DVD) and USB flash drives.
- You can attach only one USB flash drive, image, or key and one optical drive on the management station to a managed system. Supported optical drives include a maximum of one available optical drive or one ISO image file.

The following figure shows a typical Virtual Media setup.

- Any connected Virtual Media emulates a physical device on the managed system.
- On Windows-based managed systems, the Virtual Media drives are auto-mounted if they are attached and configured with a drive letter.
- On Linux-based managed systems with some configurations, the Virtual Media drives are not auto-mounted. To manually mount the drives, use the mount command.
- All the virtual drive access requests from the managed system are directed to the management station across the network.
- Virtual devices appear as two drives on the managed system without the media being installed in the drives.
- You can share the management station CD/DVD drive (read only), but not a USB media, between two managed systems.
- Virtual media requires a minimum available network bandwidth of 128 Kbps.
- If LOM or NIC failover occurs, then the Virtual Media session may be disconnected.

After attaching a Virtual Media image through Virtual Console, the drive may not show up in Windows host OS. Check Windows Device Manager for any unknown mass storage devices. Right click on the unknown device and update the driver or choose uninstall driver. The device is recognized by Windows after disconnecting and reconnecting vMedia.



Figure 4. Virtual media setup

Topics:

Inserido ao protocolo **22.951.206-4** por: **Pedro Henrique Golin Linhares** em: 31/03/2025 12:49. A autenticidade deste documento pode ser validada no endereço: https://www.eprotocolo.pr.gov.br/spiweb/validarDocumento com o código: 692a5ef82beb173ed33dad4ee6505546.



- Supported drives and devices
- Configuring virtual media
- Accessing virtual media
- Setting boot order through BIOS
- Enabling boot once for virtual media

Supported drives and devices

The following table lists the drives supported through virtual media.

Table 60. Supported drives and devices

Drive	Supported Storage Media	
Virtual Optical Drives	 CD-ROM DVD CD-RW 	
	Combination drive with CD-ROM media	
USB flash drives	USB CD-ROM drive with CD-ROM mediaUSB Key image in the ISO9660 format	

Configuring virtual media

Before you configure the Virtual Media settings, make sure that you have configured your Web browser to use Java or ActiveX plug-in.

Configuring virtual media using iDRAC web interface

To configure virtual media settings:

- CAUTION: Do not reset iDRAC when running a Virtual Media session. Otherwise, undesirable results may occur, including data loss.
- 1. In the iDRAC Web interface, go to Configuration > Virtual Media > Attached Media.
- 2. Specify the required settings. For more information, see the *iDRAC Online Help*.
- 3. Click Apply to save the settings.

Configuring virtual media using RACADM

To configure the virtual media, use the set command with the objects in the **iDRAC.VirtualMedia** group.

For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Configuring virtual media using iDRAC settings utility

You can attach, detach, or auto-attach virtual media using the iDRAC Settings utility. To do this:

- 1. In the iDRAC Settings utility, go to Media and USB Port Settings. The iDRAC Settings Media and USB Port Settings page is displayed.
- 2. In the Virtual Media section, select Detach, Attach, or Auto attach based on the requirement. For more information about the options, see *iDRAC Settings Utility Online Help*.
- **3.** Click **Back**, click **Finish**, and then click **Yes**. The Virtual Media settings are configured.



Attached media state and system response

The following table describes the system response based on the Attached Media setting.

Table 61. Attached media state and system response

Attached Media State	System Response
Detach	Cannot map an image to the system.
Attach	Media is mapped even when Client View is closed.
Auto-attach	Media is mapped when Client View is opened and unmapped when Client View is closed.

Server settings for viewing virtual devices in virtual media

You must configure the following settings in the management station to allow visibility of empty drives. To do this, in Windows Explorer, from the **Organize** menu, click **Folder and search options**. On the **View** tab, deselect **Hide empty drives in the Computer folder** option and click **OK**.

Accessing virtual media

You can access Virtual Media with or without using the Virtual Console. Before you access Virtual Media, make sure to configure your Web browser(s).

Virtual Media and RFS are mutually exclusive. If the RFS connection is active and you attempt to launch the Virtual Media client, the following error message is displayed: *Virtual Media is currently unavailable. A Virtual Media or Remote File Share session is in use.*

If the RFS connection is not active and you attempt to launch the Virtual Media client, the client launches successfully. You can then use the Virtual Media client to map devices and files to the Virtual Media virtual drives.

Launching virtual media using virtual console

Before you launch Virtual Media through the Virtual Console, make sure that:

- Virtual Console is enabled.
- System is configured to not hide empty drives In Windows Explorer, navigate to Folder Options, clear the Hide empty drives in the Computer folder option, and click OK.

To access Virtual Media using Virtual Console:

- In the iDRAC web interface, go to Configuration > Virtual Console. The Virtual Console page is displayed.
- 2. Click Launch Virtual Console.

The Virtual Console Viewer is launched.

NOTE: On Linux, Java is the default plug-in type for accessing the Virtual Console. On Windows, open the .jnlp file to launch the Virtual Console using Java.

3. Click Virtual Media > Connect Virtual Media.

The Virtual Media session is established and the Virtual Media menu displays the list of devices available for mapping.

(i) NOTE: The Virtual Console Viewer window must remain active while you access the Virtual Media.

Launching virtual media without using virtual console

Before you launch Virtual Media when the **Virtual Console** is disabled, ensure that System is configured to unhide empty drives. To do this, in Windows Explorer, go to **Folder Options**, clear the **Hide empty drives in the Computer folder** option, and click **OK**.

To access Virtual Media when Virtual Console is disabled:

304 Managing virtual media



- 1. In the iDRAC web Interface, go to **Configuration** > Virtual Media.
- 2. Click Connect Virtual Media.

Alternatively, you can also launch the Virtual Media by following these steps:

- 1. Go to Configuration > Virtual Console.
- 2. Click Launch Virtual Console. The following message is displayed:

Virtual Console has been disabled. Do you want to continue using Virtual Media redirection?

- 3. Click OK. The Virtual Media window is displayed.
- 4. From the Virtual Media menu, click Map CD/DVD or Map Removable Disk. For more information, see Mapping virtual drive.
- 5. Virtual Media Statistics shows the list of target drives, their mapping, status (Read-Only or Not), Duration of connection, Read/Write Bytes, and the transfer rate.
- (i) **NOTE:** The virtual device drive letters on the managed system do not coincide with the physical drive letters on the management station.

() NOTE: The Virtual Media may not function correctly on systems running Windows operating system configured with Internet Explorer Enhanced Security. To resolve this issue, see the Microsoft operating system documentation or contact the system administrator.

Adding virtual media images

You can create a media image of the remote folder and mount it as a USB attached device to the server's operating system. To add Virtual Media images:

- 1. Click Virtual Media > Create Image....
- 2. In the **Source Folder** field, click **Browse** and browse to the folder or directory to be used as the source for the image file. The image file is on the management station or the C: drive of the managed system.
- **3.** In the **Image File Name** field, the default path to store the created image files (typically the desktop directory) appears. To change this location, click **Browse** and navigate to a location.
- 4. Click Create Image.

The image creation process starts. If the image file location is within the source folder, a warning message is displayed indicating that the image creation cannot proceed as the image file location within the source folder causes an infinite loop. If the image file location is not within the source folder, then the image creation proceeds.

After the image is created, a success message is displayed.

5. Click Finish.

The image is created.

When a folder is added as an image, a **.img** file is created on the Desktop of the management station from which this feature is used. If this **.img** file is moved or deleted, then the corresponding entry for this folder in the **Virtual Media** menu does not work. Therefore, it is recommended not to move or delete the **.img** file while the *image* is being used. However, the **.img** file can be removed after the relevant entry is first deselected and then removed using **Remove Image** to remove the entry.

Viewing virtual device details

To view the virtual device details, in the Virtual Console Viewer, click **Tools** > **Stats**. In the **Stats** window, the **Virtual Media** section displays the mapped virtual devices and the read/write activity for each device. If Virtual Media is connected, this information is displayed. If Virtual Media is not connected, the "Virtual Media is not connected" message is displayed.

If the Virtual Media is launched without using the Virtual Console, then the **Virtual Media** section is displayed as a dialog box. It provides information about the mapped devices.

Accessing drivers

Dell EMC PowerEdge servers have all the supported operating system drivers embedded on the system flash memory. Using iDRAC, you can mount or unmount drivers easily to deploy the operating system on your server.

To mount the drivers:



- 1. On the iDRAC web interface, go to Configuration > Virtual Media.
- 2. Click Mount Drivers.
- 3. Select the OS from the pop-up window and click Mount Drivers.

(i) NOTE: The Expose duration is 18 hours by default.

To unmount the drivers post completion of the mount:

- 1. Go to Configuration > Virtual Media.
- 2. Click Unmount Drivers.
- **3.** Click \mathbf{OK} on the pop-up window.
- (i) **NOTE:** The **Mount Drivers** option may not be displayed if the driver pack is not available on the system. Ensure to download and install the latest driver pack from https://www.dell.com/support.

Resetting USB

To reset the USB device:

- In the Virtual Console viewer, click Tools > Stats. The Stats window is displayed.
- Under Virtual Media, click USB Reset.
 A message is displayed warning the user that resetting the USB connection can affect all the input to the target device including Virtual Media, keyboard, and mouse.
- 3. Click Yes.
 - The USB is reset.

(i) NOTE: iDRAC Virtual Media does not terminate even after you log out of iDRAC Web interface session.

Mapping virtual drive

To map the virtual drive:

- () NOTE: While using ActiveX or Java-based Virtual Media, you must have administrative privileges to map an operating system DVD or a USB flash drive (that is connected to the management station). To map the drives, launch IE as an administrator or add the iDRAC IP address to the list of trusted sites.
- 1. To establish a Virtual Media session, from the Virtual Media menu, click Connect Virtual Media.

For each device available for mapping from the host server, a menu item appears under the **Virtual Media** menu. The menu item is named according to the device type such as:

- Map CD/DVD
- Map Removable Disk

The Map DVD/CD option can be used for ISO files and the Map Removable Disk option can be used for images.

() NOTE:

- You cannot map physical media such USB-based drives, CD, or DVD by using the HTML5 based virtual console.
- You cannot map USB keys as virtual media disks using Virtual Console/Virtual media over a RDP session.
- You cannot map physical media with NTFS format in ehtml removable media, use FAT or exFAT devices
- 2. Click the device type that you want to map.

NOTE: The active session displays if a Virtual Media session is currently active from the current Web interface session, from another Web interface session.

3. In the Drive/Image File field, select the device from the drop-down list.

The list contains all the available (unmapped) devices that you can map (CD/DVD and Removable Disk) and image file types that you can map (ISO or IMG). The image files are located in the default image file directory (typically the user's desktop). If the device is not available in the drop-down list, click **Browse** to specify the device.

The correct file type for CD/DVD is ISO and for removable disk it is IMG.

If the image is created in the default path (Desktop), when you select **Map Removable Disk**, the created image is available for selection in the drop-down menu.



If image is created in a different location, when you select **Map Removable Disk**, the created image is not available for selection in the drop-down menu. Click **Browse** to specify the image.

() NOTE:

- **Read only** option will be grayed out in ehtml5 based JAVA removable media.
- Floppy emulation is not supported in ehtml5 plugin.
- 4. Select Read-only to map writable devices as read-only.

For CD/DVD devices, this option is enabled by default and you cannot disable it.

(i) NOTE: The ISO and IMG files map as read-only files if you map these files by using the HTML5 virtual console.

5. Click Map Device to map the device to the host server.

After the device/file is mapped, the name of its **Virtual Media** menu item changes to indicate the device name. For example, if the CD/DVD device is mapped to an image file named foo.iso, then the CD/DVD menu item on the Virtual Media menu is named **foo.iso mapped to CD/DVD**. A check mark for that menu item indicates that it is mapped.

Displaying correct virtual drives for mapping

On a Linux-based management station, the Virtual Media **Client** window may display removable disks that are not part of the management station. To make sure that the correct virtual drives are available to map, you must enable the port setting for the connected SATA hard drive. To do this:

- 1. Reboot the operating system on the management station. During POST, press <F2> to enter System Setup.
- 2. Go to SATA settings. The port details are displayed.
- 3. Enable the ports that are actually present and connected to the hard drive.
- 4. Access the Virtual Media **Client** window. It displays the correct drives that can be mapped.

Unmapping virtual drive

To unmap the virtual drive:

- 1. From the Virtual Media menu, do any of the following:
 - Click the device that you want to unmap.
 - Click Disconnect Virtual Media.

A message appears asking for confirmation.

- 2. Click Yes.
 - The check mark for that menu item does not appear indicating that it is not mapped to the host server.
 - () NOTE: After unmapping a USB device attached to vKVM from a client system running the Macintosh operating system, the unmapped device may be unavailable on the client. Restart the system or manually mount the device on the client system to view the device.

(i) NOTE: To unmap a virtual DVD drive on Linux OS, unmount the drive and eject it.

Setting boot order through BIOS

Using the System BIOS Settings utility, you can set the managed system to boot from virtual optical drives or virtual floppy drives.

(i) NOTE: Changing Virtual Media while connected may stop the system boot sequence.

To enable the managed system to boot:

- 1. Boot the managed system.
- 2. Press <F2> to enter the System Setup page.
- 3. Go to System BIOS Settings > Boot Settings > BIOS Boot Settings > Boot Sequence.

In the pop-up window, the virtual optical drives and virtual floppy drives are listed with the standard boot devices.

4. Make sure that the virtual drive is enabled and listed as the first device with bootable media. If required, follow the on-screen instructions to modify the boot order.



- 5. Click OK, navigate back to System BIOS Settings page, and click Finish.
- 6. Click Yes to save the changes and exit.

The managed system reboots.

The managed system attempts to boot from a bootable device based on the boot order. If the virtual device is connected and a bootable media is present, the system boots to the virtual device. Otherwise, the system overlooks the device—similar to a physical device without bootable media.

Enabling boot once for virtual media

You can change the boot order only once when you boot after attaching remote Virtual Media device.

- Before you enable the boot once option, make sure that:
- You have *Configure User* privilege.
- Map the local or virtual drives (CD/DVD, Floppy, or USB flash device) with the bootable media or image using the Virtual Media options
- Virtual Media is in Attached state for the virtual drives to appear in the boot sequence.
- To enable the boot once option and boot the managed system from the Virtual Media:
- 1. In the iDRAC Web interface, go to **Overview** > **Server** > **Attached Media**.
- 2. Under Virtual Media, select the Enable Boot Once and click Apply.
- 3. Turn on the managed system and press <F2> during boot.
- 4. Change the boot sequence to boot from the remote Virtual Media device.
- 5. Reboot the server.
- The managed system boots once from the Virtual Media.



Managing vFlash SD card

(i) NOTE: vFlash is supported on AMD platform servers.

The vFlash SD card is a Secure Digital (SD) card that can be ordered and installed from the factory. You can use a card with a maximum of 16 GB capacity. After you insert the card, you must enable vFlash functionality to create and manage partitions. vFlash is a licensed feature.

(i) **NOTE:** There is no limitation of the size of SD card, you can open and replace the factory installed SD card with a higher capacity SD card. Since vFlash uses FAT32 file system, file size is limited to 4GB.

If the card is not available in the system's vFlash SD card slot, the following error message is displayed in the iDRAC Web interface at **Overview** > **Server** > **vFlash**:

SD card not detected. Please insert an SD card of size 256MB or greater.

() NOTE: Make sure that you only insert a vFlash compatible SD card in the iDRAC vFlash card slot. If you insert a noncompatible SD card, the following error message is displayed when you initialize the card: *An error has occurred while initializing SD card.*

The key features are:

- Provides storage space and emulates USB device (s).
- Create up to 16 partitions. These partitions, when attached, are exposed to the system as a Floppy drive, Hard Disk drive, or a CD/DVD drive depending on the selected emulation mode.
- Create partitions from supported file system types. Supports .img format for floppy, .iso format for CD/DVD, and both .iso and .img formats for Hard Disk emulation types.
- Create bootable USB device(s).
- Boot once to an emulated USB device.

NOTE: It is possible that a vFlash license may expire during a vFlash operation. If it happens, the on-going vFlash operations complete normally.

(i) NOTE: If FIPS mode is enabled, you cannot perform any vFlash actions.

Topics:

- Configuring vFlash SD card
- Managing vFlash partitions

Configuring vFlash SD card

Before configuring vFlash, make sure that the vFlash SD card is installed on the system. For information on how to install and remove the card from your system, see the *Installation and Service Manual* available at https://www.dell.com/poweredgemanuals.

(i) NOTE: You must have Access Virtual Media privilege to enable or disable vFlash functionality, and initialize the card.

Viewing vFlash SD card properties

After vFlash functionality is enabled, you can view the SD card properties using iDRAC Web interface or RACADM.



Viewing vFlash SD card properties using web interface

To view the vFlash SD card properties, in the iDRAC Web interface, go to **Configuration** > **System Settings** > **Hardware Settings** > **vFlash**. The Card Properties page is displayed. For information about the displayed properties, see the *iDRAC Online Help*.

Viewing vFlash SD card properties using RACADM

To view the vFlash SD card properties using RACADM, use the get command with the following objects:

- iDRAC.vflashsd.AvailableSize
- iDRAC.vflashsd.Health
- iDRAC.vflashsd.Licensed
- iDRAC.vflashsd.Size
- iDRAC.vflashsd.WriteProtect

For more information about these objects, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Viewing vFlash SD card properties using iDRAC settings utility

To view the vFlash SD card properties, in the **iDRAC Settings Utility**, go to **Media and USB Port Settings**. The **Media and USB Port Settings** page displays the properties. For information about the displayed properties, see the *iDRAC Settings Utility Online Help*.

Enabling or disabling vFlash functionality

You must enable the vFlash functionality to perform partition management.

Enabling or disabling vFlash functionality using web interface

To enable or disable the vFlash functionality:

- In the iDRAC web interface, go to Configuration > System Settings > Hardware Settings > vFlash. The SD Card Properties page is displayed.
- 2. Select or clear the **vFLASH Enabled** option to enable or disable the vFlash functionality. If any vFlash partition is attached, you cannot disable vFlash and an error message is displayed.

(i) NOTE: If vFlash functionality is disabled, SD card properties are not displayed.

3. Click Apply. The vFlash functionality is enabled or disabled based on the selection.

Enabling or disabling vFlash functionality using RACADM

To enable or disable the vFlash functionality using RACADM:

racadm set iDRAC.vflashsd.Enable [n]

n=0

Disabled

n=1

Enabled

NOTE: The RACADM command functions only if a vFlash SD card is present. If a card is not present, the following message is displayed: *ERROR: SD Card not present*.



Enabling or disabling vFlash functionality using iDRAC settings utility

To enable or disable the vFlash functionality:

- In the iDRAC Settings utility, go to Media and USB Port Settings. The iDRAC Settings. Media and USB Port Settings page is displayed.
- 2. In the **vFlash Media** section, select **Enabled** to enable vFlash functionality or select **Disabled** to disable the vFlash functionality.
- Click Back, click Finish, and then click Yes. The vFlash functionality is enabled or disabled based on the selection.

Initializing vFlash SD card

The initialize operation reformats the SD card and configures the initial vFlash system information on the card.

(i) **NOTE:** If the SD card is write-protected, then the Initialize option is disabled.

Initializing vFlash SD card using web interface

To initialize the vFlash SD card:

- In the iDRAC Web interface, go to Configuration > System Settings > Hardware Settings > vFlash. The SD Card Properties page is displayed.
- 2. Enable vFLASH and click Initialize.

All existing contents are removed and the card is reformatted with the new vFlash system information.

If any vFlash partition is attached, the initialize operation fails and an error message is displayed.

Initializing vFlash SD card using RACADM

To initialize the vFlash SD card using RACADM:

racadm set iDRAC.vflashsd.Initialized 1

All existing partitions are deleted and the card is reformatted.

For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Initializing vFlash SD card using iDRAC settings utility

To initialize the vFlash SD card using iDRAC Settings utility:

- 1. In the iDRAC Settings utility, go to **Media and USB Port Settings**.
- The iDRAC Settings . Media and USB Port Settings page is displayed.
- 2. Click Initialize vFlash.
- 3. Click Yes. The initialization operation starts.
- 4. Click **Back** and navigate to the same **iDRAC Settings**. **Media and USB Port Settings** page to view the successful message.

All existing contents are removed and the card is reformatted with the new vFlash system information.

Getting the last status using RACADM

To get the status of the last initialize command sent to the vFlash SD card:

- 1. Open SSH or Serial console to the system and log in.
- 2. Enter the command: racadm vFlashsd status The status of commands sent to the SD card is displayed.
- 3. To get the last status of all the vflash partitions, use the command:racadm vflashpartition status -a



4. To get the last status of a particular partition, use command:racadm vflashpartition status -i (index)

(i) NOTE: If iDRAC is reset, the status of the last partition operation is lost.

Managing vFlash partitions

You can perform the following using the iDRAC Web interface or RACADM:

NOTE: An administrator can perform all operations on the vFlash partitions. Else, you must have **Access Virtual Media** privilege to create, delete, format, attach, detach, or copy the contents for the partition.

- Creating an empty partition
- Creating a partition using an image file
- Formatting a partition
- Viewing available partitions
- Modifying a partition
- Attaching or detaching partitions
- Deleting existing partitions
- Downloading partition contents
- Booting to a partition

() NOTE: If you click any option on the vFlash pages when an application such as WSMan, iDRAC Settings utility, or RACADM is using vFlash, or if you navigate to some other page in the GUI, iDRAC may display the message: vFlash is currently in use by another process. Try again after some time.

vFlash is capable of performing fast partition creation when there is no other on-going vFlash operation such as formatting, attaching partitions, and so on. Therefore, it is recommended to first create all partitions before performing other individual partition operations.

Creating an empty partition

An empty partition, when attached to the system, is similar to an empty USB flash drive. You can create empty partitions on a vFlash SD card. You can create partitions of type *Floppy* or *Hard Disk*. The partition type CD is supported only while creating partitions using images.

Before creating an empty partition, make sure that:

- You have Access Virtual Media privilege.
- The card is initialized.
- The card is not write-protected.
- An initialize operation is not being performed on the card.

Creating an empty partition using the web interface

To create an empty vFlash partition:

- 1. In iDRAC Web interface, go to Configuration > Systems Settings > Hardware Settings > vFlash > Create Empty Partition.
 - The Create Empty Partition page is displayed.
- Specify the required information and click Apply. For information about the options, see the *iDRAC Online Help*.
 A new unformatted empty partition is created that is read-only by default. A page indicating the progress percentage is displayed. An error message is displayed if:
 - The card is write-protected.
 - The label name matches the label of an existing partition.
 - A non-integer value is entered for the partition size, the value exceeds the available space on the card, or the partition size is greater than 4 GB.
 - An initialize operation is being performed on the card.



Creating an empty partition using RACADM

To create an empty partition:

- 1. Log in to the system using SSH or Serial console.
- 2. Enter the command:

racadm vflashpartition create -i 1 -o drive1 -t empty -e HDD -f fat16 -s [n]

where [n] is the partition size.

By default, an empty partition is created as read-write.

If the share is not configured using Username / Password, you need to specify the parameters as

```
-u anonymous -p anonymous
```

Creating a partition using an image file

You can create a new partition on the vFlash SD card using an image file (available in the **.img** or **.iso** format.) The partitions are of emulation types: Floppy (**.img**), Hard Disk (**.img**), or CD (**.iso**). The created partition size is equal to the image file size.

Before creating a partition from an image file, make sure that:

- You have Access Virtual Media privilege.
- The card is initialized.
- The card is not write-protected.
- An initialize operation is not being performed on the card.
- The image type and the emulation type match.
 - **NOTE:** The uploaded image and the emulation type must match. There are issues when iDRAC emulates a device with incorrect image type. For example, if the partition is created using an ISO image and the emulation type is specified as Hard Disk, then the BIOS cannot boot from this image.
- Image file size is less than or equal to the available space on the card.
- Image file size is less than or equal to 4 GB as the maximum partition size supported is 4 GB. However, while creating a partition using a Web browser, the image file size must be less than 2 GB.

(i) NOTE: The vFlash partition is an image file on a FAT32 file system. Thus, the image file has the 4 GB limitation.

(i) NOTE: Installation of a full OS is not supported.

Creating a partition using an image file using web interface

To create a vFlash partition from an image file:

- In iDRAC Web interface, go to Configuration > System Settings > Hardware Settings > vFlash > Create From Image. The Create Partition from Image File page is displayed.
- 2. Enter the required information and click **Apply**. For information about the options, see the *iDRAC Online Help*.

A new partition is created. For CD emulation type, a read-only partition is created. For Floppy or Hard Disk emulation type, a read-write partition is created. An error message is displayed if:

- The card is write-protected
- The label name matches the label of an existing partition.
- The size of the image file is greater than 4 GB or exceeds the available space on the card.
- The image file does not exist or the image file extension is neither .img nor .iso.
- An initialize operation is already being performed on the card.

Creating a partition from an image file using RACADM

To create a partition from an image file using RACADM:



- 1. Log in to the system using SSH or Serial console.
- 2. Enter the command

```
racadm vflashpartition create -i 1 -o drivel -e HDD -t image -l //myserver/sharedfolder/foo.iso -u root -p mypassword
```

By default, the created partition is read-only. This command is case sensitive for the image file name extension. If the file name extension is in upper case, for example FOO.ISO instead of FOO.iso, then the command returns a syntax error.

NOTE: Creating vFlash partition from an image file located on the CFS or NFS IPv6 enabled network share is not supported.

If the share is not configured using Username / Password, you need to specify the parameters as

```
-u anonymous -p anonymous
```

Formatting a partition

You can format an existing partition on the vFlash SD card based on the type of file system. The supported file system types are EXT2, EXT3, FAT16, and FAT32. You can only format partitions of type Hard Disk or Floppy, and not CD. You cannot format read-only partitions.

Before creating a partition from an image file, ensure that:

- You have Access Virtual Media privilege.
- The card is initialized.
- The card is not write-protected.
- An initialize operation is not being performed on the card.

To format vFlash partition:

- In iDRAC Web interface, go to Configuration > System Settings > Hardware Settings > vFlash > Format. The Format Partition page is displayed.
- 2. Enter the required information and click Apply.

For information about the options, see the *iDRAC Online Help*.

A warning message indicating that all the data on the partition will be erased is displayed.

3. Click OK.

The selected partition is formatted to the specified file system type. An error message is displayed if:

- The card is write-protected.
- An initialize operation is already being performed on the card.

Viewing available partitions

Make sure that the vFlash functionality is enabled to view the list of available partitions.

Viewing available partitions using web interface

To view the available vFlash partitions, in the iDRAC Web interface, go to **Configuration** > **System Settings** > **Hardware Settings** > **vFlash** > **Manage**. The **Manage Partitions** page is displayed listing the available partitions and related information for each partition. For information on the partitions, see the *iDRAC Online Help*.

Viewing available partitions using RACADM

To view the available partitions and their properties using RACADM:

1. Open a SSH or Serial console to the system and log in.



- 2. Enter the following commands:
 - To list all existing partitions and its properties: racadm vflashpartition list
 - To get the status of operation on partition 1: racadm vflashpartition status -i 1
 - To get the status of all existing partitions: racadm vflashpartition status -a

(i) NOTE: The -a option is valid only with the status action.

Modifying a partition

You can change a read-only partition to read-write or vice-versa. Before modifying the partition, make sure that:

- The vFlash functionality is enabled.
- You have Access Virtual Media privileges.

(i) NOTE: By default, a read-only partition is created.

Modifying a partition using web interface

To modify a partition:

- 1. In the iDRAC Web interface, go to Configuration > System Settings > Hardware Settings > vFlash > Manage. The Manage Partitions page is displayed.
- 2. In the Read-Only column:
 - Select the checkbox for the partition(s) and click **Apply** to change to read-only.
 - Clear the checkbox for the partition(s) and click **Apply** to change to read-write.

The partitions are changed to read-only or read-write, based on the selections.

NOTE: If the partition is of type CD, the state is read-only. You cannot change the state to read-write. If the partition is attached, the check box is grayed-out.

Modifying a partition using RACADM

To view the available partitions and their properties on the card:

- 1. Log in to the system using SSH or Serial console.
- 2. Use one of the following:
 - Using set command to change the read-write state of the partition:
 - To change a read-only partition to read-write:

racadm set iDRAC.vflashpartition.<index>.AccessType 1

• To change a read-write partition to read-only:

racadm set iDRAC.vflashpartition.<index>.AccessType 0

• Using set command to specify the Emulation type:

racadm set iDRAC.vflashpartition.<index>.EmulationType <HDD, Floppy, or CD-DVD>



Attaching or detaching partitions

When you attach one or more partitions, they are visible to the operating system and BIOS as USB mass storage devices. When you attach multiple partitions, based on the assigned index, they are listed in an ascending order in the operating system and the BIOS boot order menu.

If you detach a partition, it is not visible in the operating system and the BIOS boot order menu.

When you attach or detach a partition, the USB bus in the managed system is reset. This affects applications that are using vFlash and disconnects the iDRAC Virtual Media sessions.

Before attaching or detaching a partition, make sure that:

- The vFlash functionality is enabled.
- An initialize operation is not already being performed on the card.
- You have Access Virtual Media privileges.

Attaching or detaching partitions using web interface

To attach or detach partitions:

- In the iDRAC Web interface, go to Configuration > System Settings > Hardware Settings > vFlash > Manage. The Manage Partitions page is displayed.
- 2. In the Attached column:
 - Select the checkbox for the partition(s) and click **Apply** to attach the partition(s).
 - Clear the checkbox for the partition(s) and click **Apply** to detach the partition(s).

The partitions are attached or detached, based on the selections.

Attaching or detaching partitions using RACADM

To attach or detach partitions:

- 1. Log in to the system using SSH or Serial console.
- 2. Use the following commands:
 - To attach a partition:

```
racadm set iDRAC.vflashpartition.<index>.AttachState 1
```

• To detach a partition:

```
racadm set iDRAC.vflashpartition.<index>.AttachState 0
```

Operating system behavior for attached partitions

For Windows and Linux operating systems:

- The operating system controls and assigns the drive letters to the attached partitions.
- Read-only partitions are read-only drives in the operating system.
- The operating system must support the file system of an attached partition. Else, you cannot read or modify the contents of the partition from the operating system. For example, in a Windows environment the operating system cannot read the partition type EXT2 which is native to Linux. Also, in a Linux environment the operating system cannot read the partition type NTFS which is native to Windows.
- The vFlash partition label is different from the volume name of the file system on the emulated USB device. You can change the volume name of the emulated USB device from the operating system. However, it does not change the partition label name stored in iDRAC.

Deleting existing partitions

Before deleting existing partition(s), make sure that:


- The vFlash functionality is enabled.
- The card is not write-protected.
- The partition is not attached.
- An initialize operation is not being performed on the card.

Deleting existing partitions using web interface

To delete an existing partition:

- 1. In the iDRAC Web interface, go to Configuration > System Settings > Hardware Settings > vFlash > Manage. The Manage Partitions page is displayed.
- 2. In the **Delete** column, click the delete icon for the partition that you want to delete. A message is displayed indicating that this action permanently deletes the partition.
- **3.** Click **OK**. The partition is deleted.

Deleting existing partitions using RACADM

To delete partitions:

- 1. Open a SSH or Serial console to the system and log in.
- 2. Enter the following commands:
 - To delete a partition:

racadm vflashpartition delete -i 1

• To delete all partitions, re-initialize the vFlash SD card.

Downloading partition contents

You can download the contents of a vFlash partition in the .img or .iso format to the:

• Managed system (where iDRAC is operated from)

• Network location mapped to a management station.

Before downloading the partition contents, make sure that:

- You have Access Virtual Media privileges.
- The vFlash functionality is enabled.
- An initialize operation is not being performed on the card.
- For a read-write partition, it must not be attached.

To download the contents of the vFlash partition:

- In the iDRAC Web interface, go to Configuration > System Settings > Hardware Settings > vFlash > Download. The Download Partition page is displayed.
- 2. From the Label drop-down menu, select a partition that you want to download and click Download.
 - (i) **NOTE:** All existing partitions (except attached partitions) are displayed in the list. The first partition is selected by default.
- 3. Specify the location to save the file.

The contents of the selected partition are downloaded to the specified location.

NOTE: If only the folder location is specified, then the partition label is used as the file name, along with the extension **.iso** for CD and Hard Disk type partitions, and **.img** for Floppy and Hard Disk type partitions.

Booting to a partition

You can set an attached vFlash partition as the boot device for the next boot operation.

Before booting a partition, make sure that:



- The vFlash partition contains a bootable image (in the .img or .iso format) to boot from the device.
- The vFlash functionality is enabled.
- You have Access Virtual Media privileges.

Booting to a partition using web interface

To set the vFlash partition as a first boot device, see Booting to a partition using web interface.

NOTE: If the attached vFlash partition(s) are not listed in the **First Boot Device** drop-down menu, make sure that the BIOS is updated to the latest version.

Booting to a partition using RACADM

To set a vFlash partition as the first boot device, use the iDRAC.ServerBoot object.

For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

() NOTE: When you run this command, the vFlash partition label is automatically set to boot once

(iDRAC.ServerBoot.BootOnce is set to 1.) Boot once boots the device to the partition only once and does not keep it persistently first in the boot order.



Using SMCLP

(i) NOTE: SMCLP is only supported in iDRAC versions earlier than 4.00.00.00.

The Server Management Command Line Protocol (SMCLP) specification enables CLI-based systems management. It defines a protocol for management commands transmitted over standard character oriented streams. This protocol accesses a Common Information Model Object Manager (CIMOM) using a human-oriented command set. The SMCLP is a sub-component of the Distributed Management Task Force (DMTF) SMASH initiative to streamline systems management across multiple platforms. The SMCLP specification, along with the Managed Element Addressing Specification and numerous profiles to SMCLP mapping specifications, describes the standard verbs and targets for various management task executions.

NOTE: It is assumed that you are familiar with the Systems Management Architecture for Server Hardware (SMASH) Initiative and the Server Management Working Group (SMWG) SMCLP specifications.

The SM-CLP is a subcomponent of the Distributed Management Task Force (DMTF) SMASH initiative to streamline server management across multiple platforms. The SM-CLP specification, along with the Managed Element Addressing Specification and numerous profiles to SM-CLP mapping specifications, describes the standard verbs and targets for various management task executions.

The SMCLP is hosted from the iDRAC controller firmware and supports SSH and serial-based interfaces. The iDRAC SMCLP interface is based on the SMCLP Specification Version 1.0 provided by the DMTF organization.

NOTE: Information about the profiles, extensions, and MOFs are available at https://www.dell.com/support and all DMTF information is available at **dmtf.org/standards/profiles/**.

SM-CLP commands implement a subset of the local RACADM commands. The commands are useful for scripting since you can execute these commands from a management station command line. You can retrieve the output of commands in well-defined formats, including XML, facilitating scripting and integration with existing reporting and management tools.

Topics:

- System management capabilities using SMCLP
- Running SMCLP commands
- iDRAC SMCLP syntax
- Navigating the map address space
- Using show verb
- Usage examples

System management capabilities using SMCLP

iDRAC SMCLP enables you to:

- Manage Server Power Turn on, shut down, or reboot the system
- Manage System Event Log (SEL) Display or clear the SEL records
- View iDRAC user accounts
- View system properties

Running SMCLP commands

You can run the SMCLP commands using SSH interface. Open an SSH and log in to iDRAC as an administrator. The SMCLP prompt (admin ->) is displayed.

SMCLP prompts:

- yx1x blade servers use -\$.
- yx1x rack and tower servers use admin->.
- yx2x blade, rack, and tower servers use admin->.



where, y is an alpha-numeric character such as M (for blade servers), R (for rack servers), and T (for tower servers) and x is a number. This indicates the generation of Dell PowerEdge servers.

(i) NOTE: Scripts using -\$ can use these for yx1x systems, but starting with yx2x systems one script with admin-> can be

used for blade, rack, and tower servers.

iDRAC SMCLP syntax

The iDRAC SMCLP uses the concept of verbs and targets to provide systems management capabilities through the CLI. The verb indicates the operation to perform, and the target determines the entity (or object) that runs the operation.

The SMCLP command line syntax:

<verb> [<options>] [<target>] [<properties>]

The following table provides the verbs and its definitions.

Table 62. SMCLP verbs

Verb	Definition
cd	Navigates through the MAP using the shell
set	Sets a property to a specific value
help	Displays help for a specific target
reset	Resets the target
show	Displays the target properties, verbs, and subtargets
start	Turns on a target
stop	Shuts down a target
exit	Exits from the SMCLP shell session
version	Displays the version attributes of a target
load	Moves a binary image to a specified target address from a URL

The following table provides a list of targets.

Table 63. SMCLP targets

Target	Definitions
admin1	admin domain
admin1/profiles1	Registered profiles in iDRAC
admin1/hdwr1	Hardware
admin1/system1	Managed system target
admin1/system1/capabilities1	Managed system SMASH collection capabilities



Table 63. SMCLP targets (continued)

Target	Definitions
admin1/system1/capabilities1/elecap1	Managed system target capabilities
admin1/system1/logs1	Record Log collections target
admin1/system1/logs1/log1	System Event Log (SEL) record entry
admin1/system1/logs1/log1/record*	An individual SEL record instance on the managed system
admin1/system1/settings1	Managed system SMASH collection settings
admin1/system1/capacities1	Managed system capacities SMASH collection
admin1/system1/consoles1	Managed system consoles SMASH collection
admin1/system1/sp1	Service Processor
admin1/system1/sp1/timesvc1	Service Processor time service
admin1/system1/sp1/capabilities1	Service processor capabilities SMASH collection
admin1/system1/sp1/capabilities1/clpcap1	CLP service capabilities
admin1/system1/sp1/capabilities1/ pwrmgtcap1	Power state management service capabilities on the system
admin1/system1/sp1/capabilities1/ acctmgtcap*	Account management service capabilities
admin1/system1/sp1/capabilities1/ rolemgtcap*	Local Role Based Management capabilities
admin1/system1/sp1/capabilities1/elecap1	Authentication capabilities
admin1/system1/sp1/settings1	Service Processor settings collection
admin1/system1/sp1/settings1/clpsetting1	CLP service settings data
admin1/system1/sp1/clpsvc1	CLP service protocol service
admin1/system1/sp1/clpsvc1/clpendpt*	CLP service protocol endpoint



Table 63. SMCLP targets (continued)

Target	Definitions
admin1/system1/sp1/clpsvc1/tcpendpt*	CLP service protocol TCP endpoint
admin1/system1/sp1/jobq1	CLP service protocol job queue
admin1/system1/sp1/jobq1/job*	CLP service protocol job
admin1/system1/sp1/pwrmgtsvc1	Power state management service
admin1/system1/sp1/account1-16	Local user account
admin1/sysetm1/sp1/account1-16/identity1	Local user identity account
admin1/sysetm1/sp1/account1-16/identity2	IPMI identity (LAN) account
admin1/sysetm1/sp1/account1-16/identity3	IPMI identity (Serial) account
admin1/sysetm1/sp1/account1-16/identity4	CLP identity account
admin1/system1/sp1/acctsvc2	IPMI account management service
admin1/system1/sp1/acctsvc3	CLP account management service
admin1/system1/sp1/rolesvc1	Local Role Base Authorization (RBA) service
admin1/system1/sp1/rolesvc1/Role1-16	Local role
admin1/system1/sp1/rolesvc1/Role1-16/ privilege1	Local role privilege
admin1/system1/sp1/rolesvc2	IPMI RBA service
admin1/system1/sp1/rolesvc2/Role1-3	IPMI role
admin1/system1/sp1/rolesvc2/Role4	IPMI Serial Over LAN (SOL) role
admin1/system1/sp1/rolesvc3	CLP RBA Service
admin1/system1/sp1/rolesvc3/Role1-3	CLP role



Table 63. SMCLP targets (continued)

Target	Definitions
admin1/system1/sp1/rolesvc3/Role1-3/ privilege1	CLP role privilege

Navigating the map address space

Objects that can be managed with SM-CLP are represented by targets arranged in a hierarchical space called the Manageability Access Point (MAP) address space. An address path specifies the path from the root of the address space to an object in the address space.

The root target is represented by a slash (/) or a backslash (\). It is the default starting point when you log in to iDRAC. Navigate down from the root using the cd verb.

NOTE: The slash (/) and backslash (\) are interchangeable in SM-CLP address paths. However, a backslash at the end of a command line continues the command on the next line and is ignored when the command is parsed.

For example to navigate to the third record in the System Event Log (SEL), enter the following command:

->cd /admin1/system1/logs1/log1/record3

Enter the cd verb with no target to find your current location in the address space. The .. and . abbreviations work as they do in Windows and Linux: .. refers to the parent level and . refers to the current level.

Using show verb

To learn more about a target use the show verb. This verb displays the target's properties, sub-targets, associations, and a list of the SM-CLP verbs that are allowed at that location.

Using the -display option

The show -display option allows you to limit the output of the command to one or more of properties, targets, associations, and verbs. For example, to display just the properties and targets at the current location, use the following command:

show -display properties, targets

To list only certain properties, qualify them, as in the following command:

show -d properties=(userid,name) /admin1/system1/sp1/account1

If you only want to show one property, you can omit the parentheses.

Using the -level option

The show -level option executes show over additional levels beneath the specified target. To see all targets and properties in the address space, use the -l all option.

Using the -output option

The -output option specifies one of four formats for the output of SM-CLP verbs: text, clpcsv, keyword, and clpxml.

The default format is **text**, and is the most readable output. The **clpcsv** format is a comma-separated values format suitable for loading into a spreadsheet program. The **keyword** format outputs information as a list of keyword=value pairs one per line. The **clpxml** format is an XML document containing a **response** XML element. The DMTF has specified the **clpcsv** and **clpxml** formats and their specifications can be found on the DMTF website at **dmtf.org**.

The following example shows how to output the contents of the SEL in XML:

show -1 all -output format=clpxml /admin1/system1/logs1/log1



Usage examples

This section provides use case scenarios for SMCLP:

- Server power management
- SEL management
- Map target navigation

Server power management

The following examples show how to use SMCLP to perform power management operations on a managed system.

Type the following commands at the SMCLP command prompt:

• To switch off the server:

stop /system1
The following message is displayed:
system1 has been stopped successfully

• To switch on the server:

start /system1

The following message is displayed:

system1 has been started successfully

• To reboot the server:

reset /system1

The following message is displayed:

system1 has been reset successfully

SEL management

The following examples show how to use the SMCLP to perform SEL-related operations on the managed system. Type the following commands at the SMCLP command prompt:

• To view the SEL:

show/system1/logs1/log1 The following output is displayed: /system1/logs1/log1 Targets: Record1 Record2 Record3 Record4 Record5 Properties: InstanceID = IPMI:BMC1 SEL Log MaxNumberOfRecords = 512 CurrentNumberOfRecords = 5 Name = IPMI SEL

324 Using SMCLP



```
EnabledState = 2
OperationalState = 2
HealthState = 2
Caption = IPMI SEL
Description = IPMI SEL
ElementName = IPMI SEL
Commands:
cd
show
help
exit
version
To view the SEL record:
```

•

```
show/system1/logs1/log1
The following output is displayed:
/system1/logs1/log1/record4
Properties:
LogCreationClassName= CIM_RecordLog
CreationClassName= CIM LogRecord
LogName= IPMI SEL
RecordID= 1
MessageTimeStamp= 20050620100512.000000-000
Description= FAN 7 RPM: fan sensor, detected a failure
ElementName= IPMI SEL Record
Commands:
cd
show
help
exit
version
```

Map target navigation

The following examples show how to use the cd verb to navigate the MAP. In all examples, the initial default target is assumed to be /.

Type the following commands at the SMCLP command prompt:

• To navigate to the system target and reboot:

cd system1 reset The current default target is /.

• To navigate to the SEL target and display the log records:

```
cd system1
cd logs1/log1
```

show

• To display current target:

Using SMCLP 325



type cd .

- To move up one level: type cd ..
- To exit: exit



Deploying operating systems

You can use any of the following utilities to deploy operating systems to managed systems:

- Remote File Share
- Console

Topics:

- Deploying operating system using remote file share
- Deploying operating system using virtual media
- Deploying embedded operating system on SD card

Deploying operating system using remote file share

Before you deploy the operating system using Remote File Share (RFS), make sure that:

- Configure User and Access Virtual Media privileges for iDRAC are enabled for the user.
- Network share contains drivers and operating system bootable image file, in an industry standard format such as .img or .iso.

NOTE: While creating the image file, follow standard network-based installation procedures, and mark the deployment image as read-only to make sure that each target system boots and runs the same deployment procedure.

To deploy an operating system using RFS:

- 1. Using Remote File Share (RFS), mount the ISO or IMG image file to the managed system through NFS, CIFS, HTTP, or HTTPs.
 - **NOTE:** RFS using HTTP, basic, or digest authentication is not supported, no authentication is requored. For HTTPS, basic authentication is not supported, only digest authentication or no authentication is supported.
- 2. Go to Configuration > System Settings > Hardware Settings > First Boot Device.
- 3. Set the boot order in the First Boot Device drop-down list to select a virtual media such as floppy, CD, DVD, or ISO.
- 4. Select the **Boot Once** option to enable the managed system to reboot using the image file for the next instance only.
- 5. Click Apply.
- 6. Reboot the managed system and follow the on-screen instructions to complete the deployment.

Managing remote file shares

Using Remote File Share (RFS) feature, you can set an ISO or IMG image file on a network share and make it available to the managed server's operating system as a virtual drive by mounting it as a CD or DVD using NFS, CIFS, HTTP or HTTPs. RFS is a licensed feature.

Remote file share supports only .img and .iso image file formats. A .img file is redirected as a virtual floppy and a .iso file is redirected as a virtual CDROM.

You must have Virtual Media privileges to perform an RFS mounting.

RFS and Virtual Media features are mutually exclusive.

- If the Virtual Media client is not active, and you attempt to establish an RFS connection, the connection is established and the remote image is available to the host operating system.
- If the Virtual Media client is active, and you attempt to establish an RFS connection, the following error message is displayed:

Virtual Media is detached or redirected for the selected virtual drive.



The connection status for RFS is available in iDRAC log. Once connected, an RFS-mounted virtual drive does not disconnect even if you log out from iDRAC. The RFS connection is closed if iDRAC is reset or the network connection is dropped. The Web interface and command-line options are also available in CMCOME Modular and iDRAC to close the RFS connection. The RFS connection from CMC always overrides an existing RFS mount in iDRAC.

() NOTE:

- CIFS and NFS supports both IPv4 and IPv6 addresses.
- When the iDRAC is configured with both IPv4 and IPv6, the DNS server can contain records associating the iDRAC hostname to both addresses. If IPv4 option is disabled in iDRAC, then iDRAC may not be able to access the external IPv6 share. This is because the DNS server may still contain IPv4 records, and DNS name resolution can return the IPv4 address. In such cases, it is recommended to delete the IPv4 DNS records from the DNS server, when disabling IPv4 option in iDRAC.
- If you are using CIFS and are part of an Active Directory domain, enter the domain name with the IP address in the image file path.
- If you want to access a file from an NFS share, configure the following share permissions. These permissions are required because iDRAC interfaces run in non-root mode.
 - $\circ~$ Linux: Ensure that the share permissions are set to at least Read for the Others account.
 - Windows: Go to the **Security** tab of the share properties and add **Everyone** to **Groups or user names** field with **Read & execute** privilege.
- If ESXi is running on the managed system and if you mount a floppy image (**.img**) using RFS, the connected floppy image is not available to the ESXi operating system.
- iDRAC vFlash feature and RFS are not related.
- Only English ASCII characters are supported in network share file paths.
- The OS drive eject feature is not supported when virtual media is connected using RFS.
- RFS through HTTP or HTTPs feature is not available on CMC web interface.

Configuring remote file share using web interface

To enable remote file sharing:

- 1. In iDRAC web interface, go to **Configuration** > **Virtual Media** > **Attached Media**. The **Attached Media** page is displayed.
- 2. Under Attached Media, select Attach or Auto Attach.
- 3. Under **Remote File Share**, specify the image file path, domain name, user name, and password. For information about the fields, see the *iDRAC Online Help*.

Example for image file path:

- $\bullet~\mbox{CIFS} //<\mbox{IP}$ to connect for CIFS file system>/<file path>/<image name>
- NFS < IP to connect for NFS file system>:/<file path>/<image name>
- HTTP http://<URL>/<file path>/<image name>
- HTTPs https://<URL>/<file path>/<image name>

NOTE: To avoid I/O errors when using CIFS shares hosted on Windows 7 systems, modify the following registry keys:

- Set HKLM\SYSTEM\CurrentControlSet\Control\Session Manager\Memory Management\LargeSystemCache to 1
- Set HKLM\SYSTEM\CurrentControlSet\Services\LanmanServer\Parameters\Size to 3

i NOTE: Both '/' or '\' characters can be used for the file path.

CIFS supports both IPv4 and IPv6 addresses but NFS supports only IPv4 address.

If you are using NFS share, make sure that you provide the exact <file path> and <image name> as it is case-sensitive.

NOTE: For information on recommended characters for user names and passwords, see Recommended characters in user names and passwords.



() NOTE: The characters allowed in user names and passwords for network shares are determined by the network-share type. iDRAC supports valid characters for network share credentials as defined by the share type, except <, >, and , (comma).

4. Click Apply and then click Connect.

After the connection is established, the Connection Status displays Connected.

NOTE: Even if you have configured remote file sharing, the Web interface does not display user credential information due to security reasons.

NOTE: If the image path contains user credentials, use HTTPS to avoid credentials from displaying in the GUI and RACADM. If entering the credentials in the URL, avoid using "@" symbol, because it is a separator character.

For Linux distributions, this feature may require a manual mount command when operating at runlevel init 3. The syntax for the command is:

mount /dev/OS_specific_device / user_defined_mount_point

Where, user defined mount point is any directory you choose to use for the mount similar to any mount command.

For RHEL, the CD device (.iso virtual device) is /dev/scd0 and floppy device (.img virtual device) is /dev/sdc.

For SLES, the CD device is /dev/sr0 and the floppy device is /dev/sdc. To make sure that the correct device is used (for either SLES or RHEL), when you connect the virtual device, on the Linux OS you must immediately run the command:

tail /var/log/messages | grep SCSI

This displays the text that identifies the device (example, SCSI device sdc). This procedure also applies to Virtual Media when you are using Linux distributions in runlevel init 3. By default, the virtual media is not auto-mounted in init 3.

Configuring remote file share using RACADM

To configure remote file share using RACADM, use:

racadm remoteimage

```
racadm remoteimage <options>
```

Options are:

- -c: connect image
- -d : disconnect image
- -u <username>: username to access the network share
- -p <password>: password to access the network share

-1 <image_location>: image location on the network share; use double quotes around the location. See examples for image file path in Configuring Remote File Share Using Web Interface section

- -s : display current status
- () NOTE: All characters including alphanumeric and special characters are allowed as part of user name, password, and image_location except the following characters: ' (single quote), "(double quote), ,(comma), < (less than), and > (greater than).

() NOTE: To avoid I/O errors when using CIFS shares hosted on Windows 7 systems, modify the following registry keys:

- Set HKLM\SYSTEM\CurrentControlSet\Control\Session Manager\Memory Management\LargeSystemCache to 1
- Set HKLM\SYSTEM\CurrentControlSet\Services\LanmanServer\Parameters\Size to 3

Deploying operating system using virtual media

Before you deploy the operating system using Virtual Media, make sure that:

Virtual Media is in Attached state for the virtual drives to appear in the boot sequence.



- If Virtual Media is in Auto Attached mode, the Virtual Media application must be launched before booting the system.
- Network share contains drivers and operating system bootable image file, in an industry standard format such as .img or .iso.

To deploy an operating system using Virtual Media:

- **1.** Do one of the following:
 - Insert the operating system installation CD or DVD into the management station CD or DVD drive.
 - Attach the operating system image.
- 2. Select the drive on the management station with the required image to map it.
- 3. Use one of the following methods to boot to the required device:
 - Set the boot order to boot once from Virtual Floppy or Virtual CD/DVD/ISO using the iDRAC Web interface.
 - Set the boot order through **System Setup** > **System BIOS Settings** by pressing <F2> during boot.
- 4. Reboot the managed system and follow the on-screen instructions to complete the deployment.

Installing operating system from multiple disks

- 1. Unmap the existing CD/DVD.
- 2. Insert the next CD/DVD into the remote optical drive.
- **3.** Remap the CD/DVD drive.

Deploying embedded operating system on SD card

To install an embedded hypervisor on an SD card:

- 1. Insert the two SD cards in the Internal Dual SD Module (IDSDM) slots on the system.
- 2. Enable SD module and redundancy (if required) in BIOS.
- 3. Verify if the SD card is available on one of the drives when you <F11> during boot.
- 4. Deploy the embedded operating system and follow the operating system installation instructions.

Enabling SD module and redundancy in BIOS

To enable SD module and redundancy in BIOS:

- **1.** Press <F2> during boot.
- 2. Go to System Setup > System BIOS Settings > Integrated Devices.
- 3. Set the Internal USB Port to On. If it is set to Off, the IDSDM is not available as a boot device.
- If redundancy is not required (single SD card), set Internal SD Card Port to On and Internal SD Card Redundancy to Disabled.
- 5. If redundancy is required (two SD cards), set Internal SD Card Port to On and Internal SD Card Redundancy to Mirror.
- 6. Click **Back** and click **Finish**.
- 7. Click Yes to save the settings and press <Esc> to exit System Setup.

About IDSDM

Internal Dual SD Module (IDSDM) is available only on applicable platforms. IDSDM provides redundancy on the hypervisor SD card by using another SD card that mirrors the first SD card's content.

Either of the two SD cards can be the master. For example, if two new SD cards are installed in the IDSDM, SD1 is active (master) card and SD2 is the standby card. The data is written on both the cards, but the data is read from SD1. At any time if SD1 fails or is removed, SD2 automatically become the active (master) card.

You can view the status, health, and the availability of IDSDM using iDRAC Web Interface or RACADM. The SD card redundancy status and failure events are logged to SEL, displayed on the front panel, and PET alerts are generated if alerts are enabled.



Troubleshooting managed system using iDRAC

You can diagnose and troubleshoot a remote managed system using:

- Diagnostic console
- Post code
- Boot and crash capture videos
- Last system crash screen
- System event logs
- Lifecycle logs
- Front panel status
- Trouble indicators
- System health

Topics:

- Using diagnostic console
- Viewing post codes
- Viewing boot and crash capture videos
- Viewing logs
- Viewing last system crash screen
- Viewing System status
- Hardware trouble indicators
- Viewing system health
- Checking server status screen for error messages
- Restarting iDRAC
- Reset to Custom Defaults (RTD)
- Erasing system and user data
- Resetting iDRAC to factory default settings

Using diagnostic console

iDRAC provides a standard set of network diagnostic tools that are similar to the tools included with Microsoft Windows or Linux-based systems. Using iDRAC Web interface, you can access the network debugging tools.

To access Diagnostics Console:

- 1. In the iDRAC Web interface, go to Maintenance > Diagnostics. The Diagnostics Console Command page is displayed.
- 2. In the **Command** text box, enter a command and click **Submit**. For information about the commands, see the *iDRAC Online Help*.

The results are displayed on the same page.

Reset iDRAC and Reset iDRAC to default

- 1. In the iDRAC Web interface, go to **Maintenance** > **Diagnostics**. You have the following options:
 - Click **Reset iDRAC** to reset the iDRAC. A normal reboot operation is performed on the iDRAC. After reboot, refresh the browser to reconnect and log in to iDRAC.

Troubleshooting managed system using iDRAC 331



- Click **Reset iDRAC to Default Settings** to reset the iDRAC to the default settings. After you click **Reset iDRAC to Default Settings**, **Reset iDRAC to factory default** window is displayed. This action reset the iDRAC to the factory defaults. Chose any of the following options:
 - **a.** Preserve user and network settings.
 - b. Discard all settings and reset users to the shipping value (root/shipping value).
- c. Discard all settings and reset username and password.
- 2. A warning message is displayed. Click \mathbf{Ok} to proceed further.

Scheduling remote automated diagnostics

You can remotely invoke automated offline diagnostics on a server as a one-time event and return the results. If the diagnostics require a reboot, you can reboot immediately or stage it for a subsequent reboot or maintenance cycle (similar to updates). When diagnostics are run, the results are collected and stored in the internal iDRAC storage. You can then export the results to an NFS, CIFS, HTTP, or HTTPs network share using the diagnostics export racadm command. You can also run diagnostics using the appropriate WSMan command(s). For more information, see the WSMan documentation.

You must have iDRAC Express license to use remote automated diagnostics.

You can perform the diagnostics immediately or schedule it on a particular day and time, specify the type of diagnostics, and the type of reboot.

For the schedule, you can specify the following:

- Start time Run the diagnostic at a future day and time. If you specify TIME NOW, the diagnostic is run on the next reboot.
- End time Run the diagnostic until a date and time after the Start time. If it is not started by End time, it is marked as failed with End time expired. If you specify TIME NA, then the wait time is not applicable.

The types of diagnostic tests are:

- Express test
- Extended test
- Both in a sequence

The types of reboot are:

- Power cycle system
- Graceful shutdown (waits for operating system to turn off or for system restart)
- Forced Graceful shutdown (signals operating system to turn off and waits for 10 minutes. If the operating system does not turn off, the iDRAC power cycles the system)

Only one diagnostic job can be scheduled or run at one time. A diagnostic job can complete successfully, complete with error, or is unsuccessful. The diagnostic events including the results are recorded in Lifecycle Controller log. You can retrieve the results of the last diagnostic execution using remote RACADM or WSMan.

You can export the diagnostic results of the last completed diagnostics that were scheduled remotely to a network share such as CIFS, NFS, HTTP or HTTPS. The maximum file size is 5 MB.

You can cancel a diagnostic job when the status of the job is Unscheduled or Scheduled. If the diagnostic is running, then restart the system to cancel the job.

Before you run the remote diagnostics, make sure that:

- Lifecycle Controller is enabled.
- You have Login and Server Control privileges.

Scheduling remote automated diagnostics using RACADM

• To run the remote diagnostics and save the results on the local system, use the following command:

racadm diagnostics run -m <Mode> -r <reboot type> -s <Start Time> -e <Expiration Time>

• To export the last run remote diagnostics results, use the following command:

```
racadm diagnostics export -f <file name> -l <NFS / CIFS / HTTP / HTTPs share> -u
<username> -p <password>
```

For more information about the options, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.



Viewing post codes

Post codes are progress indicators from the system BIOS, indicating various stages of the boot sequence from power-on-reset, and allows you to diagnose any faults related to system boot-up. The **Post Codes** page displays the last system post code prior to booting the operating system.

To view the Post Codes, go to Maintenance > Troubleshooting > Post Code.

The Post Code page displays the system health indicator, a hexadecimal code, and a description of the code.

Viewing boot and crash capture videos

You can view the video recordings of:

- Last three boot cycles A boot cycle video logs the sequence of events for a boot cycle. The boot cycle videos are arranged in the order of latest to oldest.
- Last crash video A crash video logs the sequence of events leading to the failure.

This is a licensed feature.

iDRAC records fifty frames during boot time. Playback of the boot screens occur at a rate of 1 frame per second. If iDRAC is reset, the boot capture video is not available as it is stored in RAM and is deleted.

() NOTE:

- You must have Access Virtual Console or administrator privileges to playback the Boot Capture and Crash Capture videos.
- The video capture time displayed in the iDRAC GUI video player may differ from the video capture time displayed in other video players. The iDRAC GUI video player displays the time in the iDRAC time zone while all other video players display the time in the respective operating system time zones.
- () NOTE: DVC boot capture files are not videos. They are sequence of screens (at 1 particular resolution) taken during the course of the server boot. The DVC player converts these screens together to create the boot video. When you export the video from DVC (continuous snapshot and differences) to .mov (actual video) format, it is expected to use the same resolution, or a similar resolution, that the video was initially encoded with. Videos need to be exported at a similar resolution that they have been captured with.
- **NOTE:** The reason for the delay in boot capture file availability is because the boot capture buffer is not full after the host boot.

To view the Boot Capture screen, click Maintenance > Troubleshooting > Video Capture.

The Video Capture screen displays the video recordings. For more information, see the iDRAC Online Help.

NOTE: When embedded video controller is disabled and server has add-on video controller, then certain latency is expected with respect to boot capture. Hence, End of Post Messages of a video will be recorded in next capture.

Configuring video capture settings

To configure the video capture settings:

- In the iDRAC Web interface, go to Maintenance > Troubleshooting > Video Capture. The Video Capture page is displayed.
- 2. From the Video Capture Settings drop-down menu, select any of the following options:
 - **Disable** Boot capture is disabled.
 - Capture until buffer full Boot sequence is captured until the buffer size has reached.
 - Capture until end of POST Boot sequence is captured until end of POST.
- 3. Click Apply to apply the settings.



Viewing logs

You can view System Event Logs (SELs) and Lifecycle logs. For more information, see Viewing System Event Log and Viewing Lifecycle log.

Viewing last system crash screen

The last crash screen feature captures a screenshot of the most recent system crash, saves, and displays it in iDRAC. This is a licensed feature.

To view the last crash screen:

- 1. Make sure that the last system crash screen feature is enabled.
- 2. In iDRAC Web interface, go to Overview > Server > Troubleshooting > Last Crash Screen.

The Last Crash Screen page displays the last saved crash screen from the managed system.

Click **Clear** to delete the last crash screen.

i NOTE: Once iDRAC is reset or an AC power cycle event occurs, then the crash capture data is cleared.

Viewing System status

The System Status summarizes the status of the following components in the system:

- Summary
- Batteries
- Cooling
- CPUs
- Front Panel
- Intrusion
- Memory
- Network Devices
- Power Supplies
- Voltages
- Removable Flash Media
- Chassis Controller

You can view the status of the managed system:

- For rack and tower servers: LCD front panel and system ID LED status or LED front panel and system ID LED status.
- For blade servers: Only system ID LEDs.

Viewing system front panel LCD status

To view the LCD front panel status for applicable rack and tower servers, in iDRAC Web interface, go to **System** > **Overview** > **Front Panel**. The **Front Panel** page is displayed.

The **Front Panel** section displays the live feed of the messages currently being displayed on the LCD front panel. When the system is operating normally (indicated by solid blue color in the LCD front panel), both **Hide Error** and **UnHide Error** are grayed-out.

(i) NOTE: You can hide or unhide the errors only for rack and tower servers.

Based on the selection, the text box displays the current value. If you select User Defined, enter the required message in the text box. The character limit is 62. If you select None, home message is not displayed on the LCD.

To view LCD front panel status using RACADM, use the objects in the System.LCD group. For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.



Viewing system front panel LED status

To view the current system ID LED status, in iDRAC web interface, go to **System** > **Overview** > **Front Panel**. The **Front Panel** section displays the current front panel status:

- Solid blue No errors present on the managed system.
- Blinking blue Identify mode is enabled (regardless of managed system error presence).
- Solid amber Managed system is in failsafe mode.
- Blinking amber Errors present on managed system.

When the system is operating normally (indicated by blue Health icon on the LED front panel), then both **Hide Error** and **UnHide Error** is grayed-out. You can hide or unhide the errors only for rack and tower servers.

To view system ID LED status using RACADM, use the getled command.

For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Hardware trouble indicators

The hardware related problems are:

- Failure to power up
- Noisy fans
- Loss of network connectivity
- Hard drive failure
- USB media failure
- Physical damage

Based on the problem, use the following methods to correct the problem:

- Reseat the module or component and restart the system
- In case of a blade server, insert the module into a different bay in the chassis
- Replace hard drives or USB flash drives
- Reconnect or replace the power and network cables

If problem persists, see the *Installation and Service Manual* available at https://www.dell.com/poweredgemanuals for specific troubleshooting information about the hardware device.

CAUTION: You should only perform troubleshooting and simple repairs as authorized in your product documentation, or as directed by online or telephone service and support team. Damage due to servicing that is not authorized by Dell is not covered by your warranty. Read and follow the safety instructions that came with the product.

Viewing system health

You can view the status for the following components on the iDRAC, CMC and, OME-Modular Web interfaces:

- Batteries
- CPUs
- Cooling
- Intrusion
- Memory
- Power Supplies
- Removable Flash Media
- Voltages
- Miscellaneous

Click any component name in the **Server Health** section to view details about the component.



Checking server status screen for error messages

When a flashing amber LED is blinking, and a particular server has an error, the main Server Status Screen on the LCD highlights the affected server in orange. Use the LCD navigation buttons to highlight the affected server, then click the center button. Error and warning messages will be displayed on the second line. For the list of error messages displayed on the LCD panel, see the server's Owner's Manual.

Restarting iDRAC

You can perform a hard or soft iDRAC restart without turning off the server:

- Hard restart On the server, press and hold the LED button for 15 seconds.
- Soft restart Using iDRAC Web interface or RACADM.

Reset to Custom Defaults (RTD)

You can use Reset to Custom Defaults feature to upload a custom config file and RTD to the settings. The new settings are applied on top of preserving users and network settings.

Reset to Custom Defaults feature has following options:

- Upload Custom Default Settings
 - You can upload custom defaults settings file. This file can be obtained by exporting the Server Configuration Profile(SCP) in XML format (JSON format is not supported for this feature). The contents of the file can be modified by customer to add or delete the settings.
 - \circ $\,$ You can upload the SCP XML file using iDRAC GUI or RACADM interfaces.
 - The uploaded configurations are saved in the default database.
- Save current settings as custom defaults
 - This operation saves the current settings as default settings.
 - This is only supported via RACADM interface.
- Download custom default settings
 - You can download SCP XML for all the default settings.
 - This is only supported via RACADM interface.
- Initiate reset to custom defaults —
- The uploaded/saved default settings are applied.

Resetting iDRAC using iDRAC web interface

To reset iDRAC, do one of the following in the iDRAC Web interface:

- Upload Custom Defaults file:
 - Go to Configuration > Server Configuration Profile > Custom Defaults > Upload Custom Defaults
 - Upload the customized CustomConfigured.xml file from Local Share path
 - Click **Apply**. New Upload Custom Defaults Job is created.
- Reset to Custom Defaults:
 - When Upload Custom Defaults job is successful, go to Maintainance > Diagnostics, click Reset iDRAC to Factory Defalts option.
 - $\circ~$ Select Discard all settings and set to Custom default configuration.
 - Click **Continue** to initiate Reset to customs Defaults configuration.

Resetting iDRAC using RACADM

To restart iDRAC, use the **racreset** command. For more information, see the *Chassis Management Controller RACADM CLI Guide* available at https://www.dell.com/cmcmanuals.For more information, see the *OME* - *Modular for PowerEdge MX7000 Chassis RACADM CLI Guide* available at https://www.dell.com/openmanagemanuals

For Reset to default operations, use the following commands:



- Upload Custom Defaults file racadm -r <iDracIP> -u <username> -p <Password> set -f <filename> -t xml --customdefaults
- Save Current Settings as Default settings racadm -r <iDracIP> -u <username> -p <Password> set --savecustomdefaults
- Download Custom Default settings racadm -r <iDracIP> -u <username> -p <Password> get -f <filename> -t xml --customdefaults
- Reset to Custom Defaults Racadm -r <iDracIP> -u <username> -p <Password> racresetcfg -custom

Erasing system and user data

(i) NOTE: Erasing system and user data is not supported from iDRAC GUI.

You can erase system component(s) and user data for the following components:

- BIOS reset to default
- Embedded Diagnostics
- Embedded OS Driver Pack
- Lifecycle Controller Data
- iDRAC reset to default
- Overwrite hard drives that do not support Instant Secure Erase (ISE)
- Reset controller cache
- Reset vFLASH
- Erase Hard Drives, SSDs, and NVMes that support ISE
- Clear all OS applications

Before performing system erase, ensure that:

• You have iDRAC Server Control privilege.

• Lifecycle Controller is enabled.

The Lifecycle Controller Data option erases any content such as the LC Log, configuration database, rollback firmware, factory as-shipped logs, and the configuration information from the FP SPI (or management riser).

NOTE: The Lifecycle Controller log contains the information about the system erase request and any information generated when the iDRAC restarts. All previous information is removed.

You can delete individual or multiple system components using the SystemErase command:

racadm systemErase <BIOS | DIAG | DRVPACK | LCDATA | IDRAC >

where,

- bios BIOS reset to default
- diag Embedded Diagnostics
- drvpack Embedded OS Driver Pack
- Icdata Clear the Lifecycle Controller Data
- idrac iDRAC reset to default
- overwritepd Overwrite hard drives that do not support Instant Secure Erase (ISE)
- percnvcache Reset controller cache
- vflash Reset vFLASH
- secureerasepd Erase Hard Drives, SSDs, and NVMes that support ISE
- allapps Clears all OS applications

(i) NOTE: While erasing vFlash, ensure that all partitions on the vFlash card are detached before executing the operation.

() NOTE: If SEKM is enabled on the server, then disable SEKM using the racadm sekm disable command before using this command. This can avoid any storage devices being locked out which are secured by iDRAC, if SEKM settings are erased from iDRAC by executing this command.

For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Troubleshooting managed system using iDRAC 337



- **NOTE:** The Dell tech center link appears on the iDRAC GUI on Dell branded systems. If you erase system data by using WSMan command and want the link to appear again, reboot the host manually and wait for CSIOR to run.
- **NOTE:** After you run System Erase, the VDs may still appear. Run CSIOR after System Erase is completed and iDRAC is rebooted.

Resetting iDRAC to factory default settings

You can reset iDRAC to the factory default settings using the iDRAC Settings utility or the iDRAC Web interface.

Resetting iDRAC to factory default settings using iDRAC web interface

To reset iDRAC to factory default settings using the iDRAC Web interface:

- 1. Go to Maintenance > Diagnostics. The Diagnostics Console page is displayed.
- Click Reset iDRAC to Default Settings. The completion status is displayed in percentage. iDRAC reboots and is restored to factory defaults. The iDRAC IP is reset and is not accessible. You can configure the IP using the front panel or BIOS.

Resetting iDRAC to factory default settings using iDRAC settings utility

To reset iDRAC to factory default values using the iDRAC Settings utility:

- Go to Reset iDRAC configurations to defaults. The iDRAC Settings Reset iDRAC configurations to defaults page is displayed.
- 2. Click Yes.
- iDRAC reset starts.
- 3. Click **Back** and navigate to the same **Reset iDRAC configurations to defaults** page to view the success message.



SupportAssist Integration in iDRAC

SupportAssist allows you to create SupportAssist collections and utilize other SupportAssist features to monitor your system and datacenter. iDRAC provides an application interfaces for gathering platform information that enables support services to resolve platform and system problems. iDRAC helps you to generate a SupportAssist collection of the server and then export the collection to a location on the management station (local) or to a shared network location such as FTP, Trivial File Transfer Protocol (TFTP), HTTP, HTTPS, Common Internet File System (CIFS) or Network File Share (NFS). The collection is generated in the standard ZIP format. You can send this collection to technical support for troubleshooting or inventory collection. **Topics:**

- SupportAssist Registration
- Installing Service Module
- Server OS Proxy Information
- SupportAssist
- Service Request Portal
- Collection Log
- Generating SupportAssist Collection
- Settings
- Collection Settings
- Contact Information

SupportAssist Registration

To take advantage of the automated, proactive, and predictive features of SupportAssist, you must register your system with SupportAssist.

You can generate and save a collection locally or to a network, and also send to Dell EMC without registration.

NOTE: Some OEM customers do not have the Model name. The back-end Support Assist does not allow registering such systems with DELL.

Contact and shipping information

To complete the registration, you must provide the contact and shipping information.

Primary contact information

Enter Company name, Country, First Name*, Last Name*, Phone Number*, Alternate Number, and Email Address*. Check if the details are displayed correctly and make changes if you want to edit any of the fields.

*indicates that the fields are mandatory.

Secondary contact information

Enter First Name, Last name, Phone Number, Alternate Number, Email Address and check if the details are displayed correctly and make changes if you want to edit any of the fields.





Auto dispatch

When a critical event is reported to Dell-EMC through iDRAC that is registered for SupportAssist, auto dispatch workflow may be initiated. This workflow is based on the event being forwarded and registered device SupportAssist warranty level. You must enter the **Dispatch information** during the SupportAssist registration process to enable auto dispatch workflow. If onsite support is required along with dispatch parts then select **Parts Dispatch with Onsite Support**.

NOTE: Auto dispatch is enabled in systems with iDRAC Service Module (iSM) v3.4.0 for Windows. Future iSM releases will support auto dispatch for additional operating systems.

Dispatch Address

Enter an address and the preferred contact hours.

End-user license agreement

After providing all the required information, you need to accept the End User License Agreement (EULA) to complete the registration process. You have the option to print the EULA for further references. You can cancel and terminate the registration process at any point of time.

Installing Service Module

In order to register and use SupportAssist, you must have iDRAC Service Module (iSM) installed in the system. Once you **initiate Service Module Installation** you can see the installation instructions. The **Next** button remains disabled until you successfully install iSM.

Server OS Proxy Information

In case there is an issue with the connection, then the user will be prompted to provide OS proxy information. Enter **Server**, **Port**, **Username** and **Password** to configure the proxy settings.

SupportAssist

Once SupportAssist is configured, you can check the SupportAssist dash board to view the **Service Request Summary**, **Warranty Status**, **SupportAssist Overview**, **Service Requests**, and **Collection log**. Registration is not required to view or send the Collection log.

Service Request Portal

Service Request shows the Status (Open/Closed), Description, Source (Event/Phone), Service Request ID, Date Opened and Date Closed details for each event. You can select and view further details of each event. You have the option to check Service Request Portal to view additional information for any individual case.

Collection Log

Collection Log shows the details of Collection Date and Time, Collection Type (Manual, Scheduled, Event based), Data Collected (Custom Selection, All Data), Collection Status (Complete with Errors, Complete), Job ID, Sent Status, and Sent Date and Time. You can send the last persisted collection in iDRAC to Dell.

NOTE: Once generated, the Collection Log Details can be filtered to remove the Personally Identifiable Information (PII) based on the user selection.

340 SupportAssist Integration in iDRAC

Inserido ao protocolo **22.951.206-4** por: **Pedro Henrique Golin Linhares** em: 31/03/2025 12:49. A autenticidade deste documento pode ser validada no endereço: https://www.eprotocolo.pr.gov.br/spiweb/validarDocumento com o código: 692a5ef82beb173ed33dad4ee6505546.



Generating SupportAssist Collection

For generating the OS and Application logs:

- iDRAC Service Module must be installed and running in Host Operating System.
- OS Collector, which comes factory installed in iDRAC, if removed must be installed in iDRAC.

If you have to work with Tech Support on an issue with a server but the security policies restrict direct internet connection, then you can provide Tech Support with necessary data to facilitate troubleshooting of the problem without having to install software or download tools from Dell and without having access to the Internet from the server operating system or iDRAC.

You can generate a health report of the server and then export the Collection log:

- To a location on the management station (local).
- To a shared network location such as Common Internet File System (CIFS) or Network File Share (NFS). To export to a
- network share such as CIFS or NFS, direct network connectivity to the iDRAC shared or dedicated network port is required.
- To Dell EMC.

The SupportAssist Collection is generated in the standard ZIP format. The collection may contain the following information:

- Hardware inventory for all components (includes system component configuration and firmware details, Motherboard System Event Logs, iDRAC state information and Lifecycle Controller logs).
- Operating system and application information.
- Storage Controller logs.
- iDRAC Debug Logs.
- It contains an HTML5 viewer, that can be accessed once the collection is complete.
- The collection provides a massive amount of detailed system information and logs in a user friendly format that can be viewed without uploading the collection to the Tech Support site.

After the data is generated, you can view the data which contains multiple XML files and log files.

Each time the data collection is performed, an event is recorded in the Lifecycle Controller log. The event includes information such as the user who initiated the report, interface used, and the date and time of export.

On Windows, If WMI is disabled, OS Collector collection stops with an error message.

Check the appropriate privilege levels and make sure there is no firewall or security settings that may prevent from collecting the registry or software data.

Before generating the health report, make sure:

- Lifecycle Controller is enabled.
- Collect System Inventory On Reboot (CSIOR) is enabled.
- You have Login and Server Control privileges.

Generating SupportAssist Collection manually using iDRAC web interface

To generate the SupportAssist collection manually:

- 1. In the iDRAC Web interface, go to Maintenance > SupportAssist.
- 2. If the server is not registered for SupportAssist, SupportAssist Registration Wizard is displayed. Click Cancel > Cancel Registration.
- 3. Click Start a Collection.
- 4. Select the data sets to be included in the Collection.
- 5. You can opt to filter the collection for PII.
- 6. Select the destination where Collection needs to be saved.
 - a. If the server is connected to the internet and the **Send Now** option is enabled, then selecting this option transmits the Collection log to Dell EMC SupportAssist.
 - b. Save locally option allows you to save the generated Collection in the local system.
 - c. Save to Network option saves the generated Collection to user defined CIFS or NFS share location.
 i) NOTE: If Save to Network is selected, and no default location is available, the provided network details will be saved as default location for future collections. If default location already exist, then the collection will use the details specified once only.



If **Save to Network** option is selected, the user provided network details is saved as defaults (if no prior network share location have been saved) for any future collections.

- 7. Click **Collect** to proceed with Collection generation.
- 8. If prompted, accept the End User Level Agreement (EULA) to continue.
 - OS and Application Data option is grayed out and not selectable if:
 - iSM is not installed or running in Host OS, or
 - OS Collector has been removed from iDRAC, or
 - OS-BMC pass through is disabled in iDRAC, or
 - cached OS Application data is not available in iDRAC from a previous collection

Settings

This page allows you to configure the collection log settings, and if registered, you can update the contact details, enable or disable email notifications, and change the language settings.

Collection Settings

You can save the collections to a preferred network location. Use **Set Archive Directory** to set the network location. You can save the collections to a preferred network location. Use Set Archive Directory to set the network location. Enter the type of Protocol (CIFS/NFS) that you want to opt for, corresponding IP Address, Share Name, Domain Name, User Name and Password before you Test Network Connection. The Test Network Connection button will confirm a connection to the destination share.

If registered, you can opt to include Identification information while sending the data to Dell in the Collection Settings.

You can enable and schedule **Automatic Collection** options to avoid any manual intervention and keep a periodical check of the system. By default, when an event is triggered and support case is opened, SupportAssist is configured to automatically collect the system logs from the device that generated the alert and upload it to Dell. You can enable or disable Automatic Collection based on events. You can schedule the Automatic collections based on your suitable requirements. The available options are, weekly, monthly, quarterly or never. You can also configure the date and time for the scheduled periodic events. You have the option to enable or disable **ProSupport Plus Recommendation Report** while configuring the Automatic Collections.

Contact Information

This page shows the contact information details that were added during the registration of SupportAssist, and allows you to update them.



Frequently asked questions

This section lists the frequently asked questions for the following:

- System Event Log
- Network security
- Active Directory
- Single Sign On
- Smart card login
- Virtual console
- Virtual media
- vFlash SD card
- SNMP authentication
- Storage devices
- iDRAC Service Module
- RACADM
- Miscellaneous

Topics:

- System Event Log
- Custom sender email configuration for iDRAC alerts
- Network security
- Telemetry streaming
- Active Directory
- Single Sign-On
- Smart card login
- Virtual console
- Virtual media
- vFlash SD card
- SNMP authentication
- Storage devices
- GPU (Accelerators)
- iDRAC Service Module
- RACADM
- Permanently setting the default password to calvin
- Miscellaneous

System Event Log

While using iDRAC Web interface through Internet Explorer, why does SEL not save using the Save As option?

This is due to a browser setting. To resolve this:

1. In Internet Explorer, go to Tools > Internet Options > Security and select the zone you are attempting to download in.

For example, if the iDRAC device is on the local intranet, select **Local Intranet** and click **Custom level...**.

- 2. In the Security Settings window, under Downloads make sure that the following options are enabled:
 - Automatic prompting for file downloads (if this option is available)
 - File download

CAUTION: To make sure that the computer used to access iDRAC is safe, under Miscellaneous, do not enable the Launching applications and unsafe files option.



Custom sender email configuration for iDRAC alerts

Alert generated email is not from Custom sender email set on Cloud based email service.

You need to register your cloud email through this process : Support.google.com.

Network security

While accessing the iDRAC Web interface, a security warning is displayed stating that the SSL certificate issued by the Certificate Authority (CA) is not trusted.

iDRAC includes a default iDRAC server certificate to ensure network security while accessing through the Web-based interface and remote RACADM. This certificate is not issued by a trusted CA. To resolve this, upload a iDRAC server certificate issued by a trusted CA (for example, Microsoft Certificate Authority, Thawte or Verisign).

Why the DNS server not registering iDRAC?

Some DNS servers register iDRAC names that contain only up to 31 characters.

When accessing the iDRAC Web-based interface, a security warning is displayed stating that the SSL certificate hostname does not match the iDRAC hostname.

iDRAC includes a default iDRAC server certificate to ensure network security while accessing through the Web-based interface and remote RACADM. When this certificate is used, the web browser displays a security warning because the default certificate that is issued to iDRAC does not match the iDRAC hostname (for example, the IP address).

To resolve this, upload an iDRAC server certificate issued to the IP address or the iDRAC hostname. When generating the CSR (used for issuing the certificate), ensure that the common name (CN) of the CSR matches the iDRAC IP address (if certificate issued to IP) or the registered DNS iDRAC name (if certificate is issued to iDRAC registered name).

To make sure that the CSR matches the registered DNS iDRAC name:

1. In iDRAC Web interface, go to **Overview** > **iDRAC Settings** > **Network**. The **Network** page is displayed.

- 2. In the Common Settings section:
 - Select the **Register iDRAC on DNS** option.
 - In the DNS iDRAC Name field, enter the iDRAC name.
- 3. Click Apply.

Why am I unable to access iDRAC from my web browser?

This issue may occur if HTTP Strict Transport Security (HSTS) is enabled. HSTS is a web security mechanism which allows web browsers to interact using only the secure HTTPS protocol, and not HTTP.

Enable HTTPS on your browser and login to iDRAC to resolve the issue.

Why am I unable to complete operations that involve a remote CIFS share?

Import/export or any other remote file share operations that involve a CIFS share fail if they use only SMBv1. Ensure that the SMBv2 protocol is enabled on the server providing SMB/CIFS share. Refer to the Operating System documentation on how to enable the SMBv2 protocol.

Telemetry streaming

Few report data are missing while streaming telemetry reports for Rsyslog servers.

Older versions of rsyslog servers may intermittently miss few report data in some reports. You can upgrade to a newer version to avoid this issue.

Active Directory

Active Directory login failed. How to resolve this?

344 Frequently asked questions



To diagnose the problem, on the Active Directory Configuration and Management page, click Test Settings. Review the test results and fix the problem. Change the configuration and run the test until the test user passes the authorization step.

In general, check the following:

- While logging in, make sure that you use the correct user domain name and not the NetBIOS name. If you have a local iDRAC user account, log into iDRAC using the local credentials. After logging in, make sure that:
 - The Active Directory Enabled option is selected on the Active Directory Configuration and Management page.
 - The DNS setting is correct on the **iDRAC Networking configuration** page.
 - The correct Active Directory root CA certificate is uploaded to iDRAC if certificate validation was enabled.
 - The iDRAC name and iDRAC Domain name matches the Active Directory environment configuration if you are using extended schema.
 - The Group Name and Group Domain Name matches the Active Directory configuration if you are using standard schema.
 - If the user and the iDRAC object is in different domain, then do not select the User Domain from Login option. Instead select Specify a Domain option and enter the domain name where the iDRAC object resides.
- Check the domain controller SSL certificates to make sure that the iDRAC time is within the valid period of the certificate.

Active Directory login fails even if certificate validation is enabled. The test results display the following error message. Why does this occur and how to resolve this?

ERROR: Can't contact LDAP server, error:14090086:SSL routines:SSL3 GET SERVER CERTIFICATE:certificate verify failed: Please check the correct Certificate Authority (CA) certificate has been uploaded to iDRAC. Please also check if the iDRAC date is within the valid period of the certificates and if the Domain Controller Address configured in iDRAC matches the subject of the Directory Server Certificate.

If certificate validation is enabled, when iDRAC establishes the SSL connection with the directory server, iDRAC uses the uploaded CA certificate to verify the directory server certificate. The most common reasons for failing certification validation are:

- iDRAC date is not within the validity period of the server certificate or CA certificate. Check the iDRAC time and the validity • period of your certificate.
- The domain controller addresses configured in iDRAC does not match the Subject or Subject Alternative Name of the directory server certificate. If you are using an IP address, read the next question. If you are using FQDN, make sure you are using the FQDN of the domain controller and not the domain. For example, servername.example.com instead of example.com.

Certificate validation fails even if IP address is used as the domain controller address. How to resolve this?

Check the Subject or Subject Alternative Name field of your domain controller certificate. Normally, Active Directory uses the host name and not the IP address of the domain controller in the Subject or Subject Alternative Name field of the domain controller certificate. To resolve this, do any of the following:

- Configure the host name (FQDN) of the domain controller as the domain controller address(es) on iDRAC to match the Subject or Subject Alternative Name of the server certificate.
- Reissue the server certificate to use an IP address in the Subject or Subject Alternative Name field, so that it matches the IP address configured in iDRAC.
- Disable certificate validation if you choose to trust this domain controller without certificate validation during the SSL handshake.

How to configure the domain controller address(es) when using extended schema in a multiple domain environment?

This must be the host name (FQDN) or the IP address of the domain controller(s) that serves the domain in which the iDRAC object resides.

When to configure Global Catalog Address(es)?

If you are using standard schema and the users and role groups are from different domains, Global Catalog Address(es) are required. In this case, you can use only Universal Group.

If you are using standard schema and all the users and role groups are in the same domain, Global Catalog Address(es) are not required.

If you are using extended schema, the Global Catalog Address is not used.

How does standard schema query work?

iDRAC connects to the configured domain controller address(es) first. If the user and role groups are in that domain, the privileges are saved.



If Global Controller Address(es) is configured, iDRAC continues to query the Global Catalog. If additional privileges are retrieved from the Global Catalog, these privileges are accumulated.

Does iDRAC always use LDAP over SSL?

Yes. All the transportation is over secure port 636 and/or 3269. During test setting, iDRAC does a LDAP CONNECT only to isolate the problem, but it does not do an LDAP BIND on an insecure connection.

Why does iDRAC enable certificate validation by default?

iDRAC enforces strong security to ensure the identity of the domain controller that iDRAC connects to. Without certificate validation, a hacker can spoof a domain controller and hijack the SSL connection. If you choose to trust all the domain controllers in your security boundary without certificate validation, you can disable it through the Web interface or RACADM.

Does iDRAC support the NetBIOS name?

Not in this release.

Why does it take up to four minutes to log in to iDRAC using Active Directory Single Sign-On or Smart Card Login?

The Active Directory Single Sign–On or Smart Card log in normally takes less than 10 seconds, but it may take up to four minutes to log in if you have specified the preferred DNS server and the alternate DNS server, and the preferred DNS server has failed. DNS time-outs are expected when a DNS server is down. iDRAC logs you in using the alternate DNS server.

The Active Directory is configured for a domain present in Windows Server 2008 Active Directory. A child or sub domain is present for the domain, the user and group is present in the same child domain, and the user is a member of that group. When trying to log in to iDRAC using the user present in the child domain, Active Directory Single Sign-On login fails.

This may be because of the an incorrect group type. There are two kinds of Group types in the Active Directory server:

- Security Security groups allow you to manage user and computer access to shared resources and to filter group policy settings.
- Distribution Distribution groups are intended to be used only as email distribution lists.

Always make sure that the group type is Security. You cannot use distribution groups to assign permission on any object, however use them to filter group policy settings.

Single Sign-On

SSO login fails on Windows Server 2008 R2 x64. What are the settings required to resolve this?

1. Run the technet.microsoft.com/en-us/library/dd560670(WS.10).aspx for the domain controller and domain policy.

2. Configure the computers to use the DES-CBC-MD5 cipher suite.

These settings may affect compatibility with client computers or services and applications in your environment. The Configure encryption types allowed for Kerberos policy setting is located at **Computer Configuration** > **Security Settings** > **Local Policies** > **Security Options**.

- 3. Make sure that the domain clients have the updated GPO.
- 4. At the command line, type gpupdate /force and delete the old key tab with klist purge command.
- 5. After the GPO is updated, create the new keytab.
- 6. Upload the keytab to iDRAC.

You can now log in to iDRAC using SSO.

Why does SSO login fail with Active Directory users on Windows 7 and Windows Server 2008 R2?

You must enable the encryption types for Windows 7 and Windows Server 2008 R2. To enable the encryption types:

- 1. Log in as administrator or as a user with administrative privilege.
- 2. Go to Start and run gpedit.msc. The Local Group Policy Editor window is displayed.
- 3. Go to Local Computer Settings > Windows Settings > Security Settings > Local Policies > Security Options.
- 4. Right-click Network Security: Configure encryption types allowed for kerberos and select Properties.
- 5. Enable all the options.

6. Click OK. You can now log in to iDRAC using SSO.

Perform the following additional settings for Extended Schema:

- 1. In the Local Group Policy Editor window, navigate to Local Computer Settings > Windows Settings > Security Settings > Local Policies > Security Options .
- 2. Right-click Network Security: Restrict NTLM: Outgoing NTLM traffic to remote server and select Properties.



- 3. Select Allow all, click OK, and close the Local Group Policy Editor window.
- 4. Go to Start and run cmd. The command prompt window is displayed.
- 5. Run the command gpupdate /force. The group policies are updated. Close the command prompt window.
- 6. Go to Start and run regedit. The Registry Editor window is displayed.
- 7. Navigate to $HKEY_LOCAL_MACHINE > System > CurrentControlSet > Control > LSA$.
- 8. In the right-pane, right-click and select New > DWORD (32-bit) Value.
- 9. Name the new key as **SuppressExtendedProtection**.
- 10. Right-click SuppressExtendedProtection and click Modify.
- 11. In the Value data field, type 1 and click OK.
- 12. Close the Registry Editor window. You can now log in to iDRAC using SSO.

If you have enabled SSO for iDRAC and you are using Internet Explorer to log in to iDRAC, SSO fails and you are prompted to enter your user name and password. How to resolve this?

Make sure that the iDRAC IP address is listed in the **Tools** > **Internet Options** > **Security** > **Trusted sites**. If it is not listed, SSO fails and you are prompted to enter your user name and password. Click **Cancel** and proceed.

Smart card login

It takes up to four minutes to log into iDRAC using Active Directory Smart Card login.

The normal Active Directory Smart Card login normally takes less than 10 seconds, however it may take up to four minutes if you have specified the preferred DNS server and the alternate DNS server in the **Network** page, and the preferred DNS server has failed. DNS time-outs are expected when a DNS server is down. iDRAC logs you in using the alternate DNS server.

ActiveX plug-in unable to detect the Smart Card reader.

Make sure that the smart card is supported on the Microsoft Windows operating system. Windows supports a limited number of smart card Cryptographic Service Providers (CSPs).

In general, check if the smart card CSPs are present on a particular client, insert the smart card in the reader at the Windows logon (Ctrl-Alt-Del) screen and check if Windows detects the smart card and displays the PIN dialog-box.

Incorrect Smart Card PIN.

Check if the smart card is locked due to too many attempts with an incorrect PIN. In such cases, contact the smart card issuer in the organization to get a new smart card.

Virtual console

What is the required Java version to launch Virtual Console?

You need Java 8 or later to use this feature and to launch iDRAC Virtual Console over an IPv6 network.

Virtual Console session is active even if you have logged out of iDRAC web interface. Is this the expected behavior?

Yes. Close the Virtual Console Viewer window to log out of the corresponding session.

Can a new remote console video session be started when the local video on the server is turned off?

Yes.

Why does it take 15 seconds to turn off the local video on the server after requesting to turn off the local video?

It gives a local user an opportunity to take any action before the video is switched off.

Is there a time delay when turning on the local video?

No, after a local video turn ON request is received by iDRAC, the video is turned on instantly.

Can the local user also turn off or turn on the video?

When the local console is disabled, the local user cannot turn off or turn on the video.

Does switching off the local video also switch off the local keyboard and mouse?

No.

Does turning off the local console turn off the video on the remote console session?

No, turning the local video on or off is independent of the remote console session.



What privileges are required for an iDRAC user to turn on or turn off the local server video?

Any user with iDRAC configuration privileges can turn on or turn off the local console.

How to get the current status of the local server video?

The status is displayed on the Virtual Console page.

To display the status of the object iDRAC.VirtualConsole.AttachState, use the following command:

racadm get idrac.virtualconsole.attachstate

Or, use the following command from a SSH or a remote session:

racadm -r (iDrac IP) -u (username) -p (password) get iDRAC.VirtualConsole.AttachState

The status is also seen on the Virtual Console OSCAR display. When the local console is enabled, a green status is displayed next to the server name. When disabled, a yellow dot indicates that iDRAC has locked the local console.

Why is the bottom of the system screen not seen from the Virtual Console window?

Make sure that the management station's monitor resolution is set to 1280 x 1024.

Why is the Virtual Console Viewer window garbled on Linux operating system?

The console viewer on Linux requires a UTF-8 character set. Check your locale and reset the character set if required.

Why does the mouse not synchronize under the Linux text console in Lifecycle Controller?

Virtual Console requires the USB mouse driver, but the USB mouse driver is available only under the X-Window operating system. In the Virtual Console viewer, do any of the following:

- Go to Tools > Session Options > Mouse tab. Under Mouse Acceleration, select Linux.
- Under the Tools menu, select Single Cursor option.

How to synchronize the mouse pointers on the Virtual Console Viewer window?

Before starting a Virtual Console session, make sure that the correct mouse is selected for your operating system.

Make sure that the **Single Cursor** option under **Tools** in the iDRAC Virtual Console menu is selected on iDRAC Virtual Console client. The default is two cursor mode.

Can a keyboard or mouse be used while installing a Microsoft operating system remotely through the Virtual Console?

No. When you remotely install a supported Microsoft operating system on a system with Virtual Console enabled in the BIOS, an EMS Connection Message is sent that requires that you select **OK** remotely. You must either select **OK** on the local system or restart the remotely managed server, reinstall, and then turn off the Virtual Console in BIOS.

This message is generated by Microsoft to alert the user that Virtual Console is enabled. To make sure that this message does not appear, always turn off Virtual Console in the iDRAC Settings utility before remotely installing an operating system.

Why does the Num Lock indicator on the management station not reflect the status of the Num Lock on the remote server?

When accessed through the iDRAC, the Num Lock indicator on the management station does not necessarily coincide with the state of the Num Lock on the remote server. The state of the Num Lock depends the setting on the remote server when the remote session is connected, regardless of the state of the Num Lock on the management station.

Why do multiple Session Viewer windows appear when a Virtual Console session is established from the local host?

You are configuring a Virtual Console session from the local system. This is not supported.

If a Virtual Console session is in-progress and a local user accesses the managed server, does the first user receive a warning message?

No. If a local user accesses the system, both have control of the system.

How much bandwidth is required to run a Virtual Console session?

It is recommended to have a 5 MBPS connection for good performance. A 1 MBPS connection is required for minimal performance.

What is the minimum system requirements for the management station to run Virtual Console?

The management station requires an Intel Pentium III 500 MHz processor with at least 256 MB of RAM.

Why does Virtual Console Viewer window sometimes displays No Signal message?



You may see this message because the iDRAC Virtual Console plug-in is not receiving the remote server desktop video. Generally, this behavior may occur when the remote server is turned off. Occasionally, the message may be displayed due to a remote server desktop video reception malfunction.

Why does Virtual Console Viewer window sometimes display an Out of Range message?

You may see this message because a parameter necessary to capture video is beyond the range for which the iDRAC can capture the video. Parameters such as display resolution or refresh rate too high causes an out of range condition. Normally, physical limitations such as video memory size or bandwidth sets the maximum range of parameters.

When starting a Virtual Console session from iDRAC web interface, why is an ActiveX security popup displayed?

iDRAC may not be in the trusted site list. To prevent the security popup from appearing every time you begin a Virtual Console session, add iDRAC to the trusted site list in the client browser:

- 1. Click Tools > Internet Options > Security > Trusted sites.
- 2. Click **Sites** and enter the IP address or the DNS name of iDRAC
- 3. Click Add.
- 4. Click Custom Level.
- 5. In the Security Settings window, select Prompt under Download unsigned ActiveX Controls.

Why is the Virtual Console Viewer window blank?

If you have Virtual Media privilege, but not Virtual Console privilege, you can start the viewer to access the virtual media feature, but the managed server's console is not displayed.

Why doesn't the mouse synchronize in DOS when using Virtual Console?

The Dell BIOS is emulating the mouse driver as a PS/2 mouse. By design, the PS/2 mouse uses relative position for the mouse pointer, which causes the lag in syncing. iDRAC has a USB mouse driver that allows absolute position and closer tracking of the mouse pointer. Even if iDRAC passes the USB absolute mouse position to the Dell BIOS, the BIOS emulation converts it back to relative position and the behavior remains. To fix this problem, set the mouse mode to USC/Diags in the Configuration screen.

After launching the Virtual Console, the mouse cursor is active on the Virtual Console, but not on the local system. Why does this occur and how to resolve this?

This occurs if the **Mouse Mode** is set to **USC/Diags**. Press **Alt + M** hot key to use the mouse on the local system. Press **Alt + M** again to use the mouse on the Virtual Console.

Why does GUI session time out after launching a virtual console from the iDRAC interface that is launched from CMC?

When launching the Virtual Console to iDRAC from the CMC web interface a popup is opened to launch the Virtual Console. The popup closes shortly after the Virtual Console opens.

When launching both the GUI and Virtual Console to the same iDRAC system on a management station, a session time-out for the iDRAC GUI occurs if the GUI is launched before the popup closes. If the iDRAC GUI is launched from the CMC web interface after the popup with the Virtual Console closed, this issue does not appear.

(i) NOTE: Not applicable for MX platforms.

Why does Linux SysRq key not work with Internet Explorer?

The Linux SysRq key behavior is different when using Virtual Console from Internet Explorer. To send the SysRq key, press the **Print Screen** key and release while holding the **Ctrl** and **Alt** keys. To send the SysRq key to a remote Linux server though iDRAC, while using Internet Explorer:

1. Activate the magic key function on the remote Linux server. You can use the following command to activate it on the Linux terminal:

echo 1 > /proc/sys/kernel/sysrq

- 2. Activate the keyboard pass-through mode of Active X Viewer.
- 3. Press Ctrl+Alt+Print Screen.
- 4. Release only Print Screen.
- 5. Press Print Screen+Ctrl+Alt.

(i) NOTE: The SysRq feature is currently not supported with Internet Explorer and Java.

Why is the "Link Interrupted" message displayed at the bottom of the Virtual Console?

When using the shared network port during a server reboot, iDRAC is disconnected while BIOS is resetting the network card. This duration is longer on 10 Gb cards, and is also exceptionally long if the connected network switch has Spanning Tree



Protocol (STP) enabled. In this case, it is recommended to enable "portfast" for the switch port connected to the server. In most cases, the Virtual Console restores itself.

Launching Virtual Console with Java plug-in fails after the iDRAC firmware was updated.

Delete the Java cache and then launch the virtual console.

To enable console redirection using the web server port (443)

racadm>>set iDRAC.VirtualConsole.WebRedirect Enabled

To close the external virtual console port (5900), set the following iDRAC property.

To close the external virtual console port (5900), both iDRAC.VirtualConsole.WebRedirect and iDRAC.VirtualConsole.CloseUnusedPort must be enabled.

racadm>>set iDRAC.VirtualConsole.CloseUnusedPort Enabled

(i) NOTE:

- If the Virtual Media port is disabled, the stand-alone Virtual Media will not be accessible and you may use the Virtual Media through Virtual Console.
- While CloseUnusedPort is enabled, Java and ActiveX based Virtual console and Virtual media will not function as they
 require dedicated external port. Virtual console and Virtual media using HTML5 plug-in will function on iDRAC web
 server port (443).

Virtual media

Why does the Virtual Media client connection sometimes drop?

When a network time-out occurs, iDRAC firmware drops the connection, disconnecting the link between the server and the virtual drive.

If you change the CD in the client system, the new CD may have an autostart feature. In this case, the firmware can time out and the connection is lost if the client system takes too long to read the CD. If a connection is lost, reconnect from the GUI and continue the previous operation.

If the Virtual Media configuration settings are changed in the iDRAC web interface or through local RACADM commands, any connected media is disconnected when the configuration change is applied.

To reconnect to the Virtual Drive, use the Virtual Media Client View window.

Why does a Windows operating system installation through Virtual Media take an extended amount of time?

If you are installing the Windows operating system using the *Dell Systems Management Tools and Documentation DVD* and the network connection is slow, the installation procedure may require an extended amount of time to access iDRAC web interface due to network latency. The installation window does not indicate the installation progress.

How to configure the virtual device as a bootable device?

On the managed system, access BIOS Setup and go to the boot menu. Locate the virtual CD, virtual floppy, or vFlash and change the device boot order as required. Also, press the "spacebar" key in the boot sequence in the CMOS setup to make the virtual device bootable. For example, to boot from a CD drive, configure the CD drive as the first device in the boot order.

What are the types of media that can be set as a bootable device?

iDRAC allows you to boot from the following bootable media:

- CDROM/DVD Data media
- ISO 9660 image
- 1.44 Floppy disk or floppy image
- A USB key that is recognized by the operating system as a removable disk
- A USB key image

How to make the USB key a bootable device?

You can also boot with a Windows 98 startup disk and copy system files from the startup disk to the USB key. For example, from the DOS prompt, type the following command:

sys a: x: /s



where, x: is the USB key that is required to be set as a bootable device.

The Virtual Media is attached and connected to the remote floppy. But, cannot locate the Virtual Floppy/Virtual CD device on a system running Red Hat Enterprise Linux or the SUSE Linux operating system. How to resolve this?

Some Linux versions do not auto-mount the virtual floppy drive and the virtual CD drive in the same method. To mount the virtual floppy drive, locate the device node that Linux assigns to the virtual floppy drive. To mount the virtual floppy drive:

1. Open a Linux command prompt and run the following command:

grep "Virtual Floppy" /var/log/messages

- 2. Locate the last entry to that message and note the time.
- 3. At the Linux prompt, run the following command:

grep "hh:mm:ss" /var/log/messages

where, hh:mm:ss is the time stamp of the message returned by grep in step 1.

- 4. In step 3, read the result of the grep command and locate the device name that is given to the Virtual Floppy.
- 5. Make sure that you are attached and connected to the virtual floppy drive.
- 6. At the Linux prompt, run the following command:

mount /dev/sdx /mnt/floppy

where, /dev/sdx is the device name found in step 4 and /mnt/floppy is the mount point.

- To mount the virtual CD drive, locate the device node that Linux assigns to the virtual CD drive. To mount the virtual CD drive:
- 1. Open a Linux command prompt and run the following command:

grep "Virtual CD" /var/log/messages

- 2. Locate the last entry to that message and note the time.
- 3. At the Linux prompt, run the following command:

grep "hh:mm:ss" /var/log/messages

where, hh:mm:ss is the timestamp of the message returned by grep in step 1.

- 4. In step 3, read the result of the grep command and locate the device name that is given to the Dell Virtual CD.
- **5.** Make sure that the Virtual CD Drive is attached and connected.
- 6. At the Linux prompt, run the following command:

mount /dev/sdx /mnt/CD

where: /dev/sdx is the device name found in step 4 and /mnt/floppy is the mount point.

Why are the virtual drives attached to the server removed after performing a remote firmware update using the iDRAC web interface?

Firmware updates cause the iDRAC to reset, drop the remote connection, and unmount the virtual drives. The drives reappear when iDRAC reset is complete.

Why are all the USB devices detached after connecting a USB device?

Virtual media devices and vFlash devices are connected as a composite USB device to the Host USB BUS, and they share a common USB port. Whenever any virtual media or vFlash USB device is connected to or disconnected from the host USB bus, all the Virtual Media and vFlash devices are disconnected momentarily from the host USB bus, and then they are reconnected. If the host operating system uses a virtual media device, do not attach or detach one or more virtual media or vFlash devices. It is recommended that you connect all the required USB devices first before using them.

What does the USB Reset do?

It resets the remote and local USB devices connected to the server.

How to maximize Virtual Media performance?

To maximize Virtual Media performance, launch the Virtual Media with the Virtual Console disabled or do one of the following:

- Change the performance slider to Maximum Speed.
- Disable encryption for both Virtual Media and Virtual Console.



NOTE: In this case, the data transfer between managed server and iDRAC for Virtual Media and Virtual Console will not be secured.

If you are using any Windows server operating systems, stop the Windows service named Windows Event Collector. To do
this, go to Start > Administrative Tools > Services. Right-click Windows Event Collector and click Stop.

While viewing the contents of a floppy drive or USB key, a connection failure message is displayed if the same drive is attached through the virtual media?

Simultaneous access to virtual floppy drives is not allowed. Close the application used to view the drive contents before attempting to virtualize the drive.

What file system types are supported on the Virtual Floppy Drive?

The virtual floppy drive supports FAT16 or FAT32 file systems.

Why is an error message displayed when trying to connect a DVD/USB through virtual media even though the virtual media is currently not in use?

The error message is displayed if Remote File Share (RFS) feature is also in use. At a time, you can use RFS or Virtual Media and not both.

Virtual Media inaccessible even though iDRAC shows Virtual Media Connection Status as Connected.

If you try to access Virtual Media using ActiveX or Java plug-in while **Attach Mode** is set to **Detach** in iDRAC, then the connection status may show as **Connected**. Change the **Attach Mode** to either **Auto-attach** or **Attach** to access the Virtual Media.

vFlash SD card

When is the vFlash SD card locked?

The vFlash SD card is locked when an operation is in-progress. For example, during an initialize operation.

SNMP authentication

Why is the message 'Remote Access: SNMP Authentication Failure' displayed?

As part of discovery, IT Assistant attempts to verify the get and set community names of the device. In IT Assistant, you have the get community name = public and the set community name = private. By default, the SNMP agent community name for iDRAC agent is public. When IT Assistant sends out a set request, the iDRAC agent generates the SNMP authentication error because it accepts requests only from community = public.

To prevent SNMP authentication errors from being generated, you must enter community names that are accepted by the agent. Since the iDRAC only allows one community name, you must use the same get and set community name for IT Assistant discovery setup.

Storage devices

OpenManage Storage Management displays more storage devices than iDRAC while information for all the storage devices connected to the system are not displayed. Why?

iDRAC displays information for only the Comprehensive Embedded Management (CEM) supported devices.

For External JBODs/Insights behind the HBA, EEMI Message for the SAS Connector/IOM removal is generated with the EEMI message ID ENC42, however, EEMI message ENC41 for the SAS Connector/IOM restoration is not generated.

To confirm restoration of the IOM in iDRAC web interface:

- 1. Go to Storage > Overview > Enclosures
- 2. Select the enclosure.
- 3. Under Advanced Properties, ensure that the value for Redundant Path is set to Present, then IOM restoration is confirmed.


GPU (Accelerators)

Accelerators section under CPU/Accelerators in iDRAC GUI is grayed out.

Few pages in GUI may not show expected response when respective attribute is disabled in Redfish.

iDRAC Service Module

iSM details are missing / not updated correctly in iDRAC GUI page of some PowerEdge servers

When a user adds SUB NIC under teaming, the configuration is invalid. This causes iSM to not to communicate with iDRAC properly.

Before installing or running the iDRAC Service Module, should the OpenManage Server Administrator be uninstalled?

No you do not have to uninstall Server Administrator. Before you install or run the iDRAC Service Module, make sure that you have stopped the features of Server Administrator that the iDRAC Service Module provides.

How to check whether iDRAC Service Module is installed in the host operating system?

To know if the iDRAC Service Module is installed on the system,

On systems running Windows:

Open the Control Panel, verify if iDRAC Service Module is listed in the list of installed programs displayed.

• On systems running Linux:

Run the command rpm -qi dcism. If the iDRAC Service Module is installed, the status displayed is **installed**.

- On systems running ESXi: Run the command esxcli software vib list|grep -i open on the host. iDRAC Service module is displayed.
- **NOTE:** To check if the iDRAC Service Module is installed on Red Hat Enterprise Linux 7, use the systemctl status dcismeng.service command instead of the init.d command.

How to check the version number of the iDRAC Service Module installed in the system?

To check the version of the iDRAC Service Module in the system, do any of the following:

- Click Start > Control Panel > Programs and Features. The version of the installed iDRAC Service Module is listed in the Version tab.
- Go to My Computer > Uninstall or change a program.

What is the minimum permission level required to install the iDRAC Service Module?

To install the iDRAC Service Module, you must have administrator level privileges.

On iDRAC Service Module version 2.0 and earlier, while installing the iDRAC Service Module, an error message is displayed stating this is not a supported server. Consult the User Guide for additional information about the supported servers. How to resolve this error?

Before installing the iDRAC Service Module, make sure that the server is a 12th generation PowerEdge server or later. Also, make sure that you have a 64-bit system.

The following message is displayed in the OS log, even when the OS to iDRAC Pass-through over USBNIC is configured properly. Why?

The iDRAC Service Module is unable to communicate with iDRAC using the OS to iDRAC Pass-through channel

iDRAC Service Module uses the OS to iDRAC pass-through over USB NIC feature to establish the communication with iDRAC. Sometimes, the communication is not established though the USB NIC interface is configured with the correct IP endpoints. This may happen when the host operating system routing table has multiple entries for the same destination mask and the USB NIC destination is not listed as the first one in routing order.

Table 64. Example of a routing order

Destination	Gateway	Genmask	Flags	Metric	Ref	Use Iface
default	10.94.148.1	0.0.0.0	UG	1024	0	0 em1
10.94.148.0	0.0.0.0	255.255.255.0	U	0	0	0 em1



Table 64. Example of a routing order (continued)

Destination	Gateway	Genmask	Flags	Metric	Ref	Use Iface
link-local	0.0.0.0	255.255.255.0	U	0	0	0 em1
link-local	0.0.0.0	255.255.255.0	U	0	0	0 enp0s20u12u3

In the example **enp0s20u12u3** is the USB NIC interface. The link-local destination mask is repeated and the USB NIC is not the first one in the order. This results in the connectivity issue between iDRAC Service Module and iDRAC over the OS to iDRAC Pass-through. To troubleshoot the connectivity issue, make sure that the iDRAC USBNIC IPv4 address (by default it is 169.254.1.1) is reachable from the host operating system.

If not:

- Change the iDRAC USBNIC address on a unique destination mask.
- Delete the entries that are not required from the routing table to make sure that USB NIC is chosen by route when the host wants to reach the iDRAC USB NIC IPv4 address.

On iDRAC Service Module version 2.0 and earlier, when uninstalling iDRAC Service Module from a VMware ESXi server, the virtual switch is named as vSwitchiDRACvusb and port group as iDRAC Network on the vSphere client. How to delete them?

While installing iDRAC Service Module VIB on a VMware ESXi server, iDRAC Service Module creates the vSwtich and Portgroup to communicate with iDRAC over the OS to iDRAC Pass-through in USB NIC mode. After the uninstallation, the virtual switch **vSwitchiDRACvusb** and the port group **iDRAC Network** are not deleted. To delete it manually, perform one of the following steps:

- Go to vSphere Client Configuration wizard and delete the entries.
- Go to the Esxcli and type the following commands:
 - To remove port group: esxcfg-vmknic -d -p "iDRAC Network"
 - To remove vSwitch: esxcfg-vswitch -d vSwitchiDRACvusb
 - **NOTE:** You can reinstall iDRAC Service Module on the VMware ESXi server as this is not a functional issue for the server.

Where is the Replicated Lifecycle log available on the operating system?

To view the replicated Lifecycle logs:

Table 65. Lifecycle logs location

Operating System	Location		
Microsoft Windows	 Event viewer > Windows Logs > System. All the iDRAC Service Module Lifecycle logs are replicated under the source name iDRAC Service Module. (i) NOTE: In iSM version 2.1 and later, Lifecycle logs are replicated under the Lifecycle Controller Log source name. In iSM version 2.0 and earlier, the logs are replicated under iDRAC Service Module source name. (i) NOTE: The location of the Lifecycle log can be configured using the iDRAC Service Module installer. You can configure the location while installer. 		
Red Hat Enterprise Linux, SUSE Linux, CentOS, and Citrix XenServer	/var/log/messages		
VMware ESXi	/var/log/syslog.log		

What are the Linux-dependent packages or executables available for installation while completing the Linux installation?

To see the list of Linux-dependent packages, see the *Linux Dependencies* section in the *iDRAC Service Module User's Guide* available at https://www.dell.com/idracmanuals.

How to increase GPU performance for certain configuration?

BIOS system performance profile set to performance



Under Processor Settings, set NPS to 4 and CCX to auto Minimum 1 DIMM per channel IOmmu=passthrough on Linux OS

RACADM

After performing an iDRAC reset (using the racadm racreset command), if any command is issued, the following message is displayed. What does this indicate?

ERROR: Unable to connect to RAC at specified IP address

The message indicates that you must wait until the iDRAC completes the reset before issuing another command.

When using RACADM commands and subcommands, some errors are not clear.

- You may see one or more of the following errors when using the RACADM commands:
- Local RACADM error messages Problems such as syntax, typographical errors, and incorrect names.
- Remote RACADM error messages Problems such as incorrect IP Address, incorrect user name, or incorrect password.

During a ping test to iDRAC, if the network mode is switched between Dedicated and Shared modes, there is no ping response.

Clear the ARP table on your system.

Remote RACADM fails to connect to iDRAC from SUSE Linux Enterprise Server (SLES) 11 SP1.

Make sure that the official openssl and libopenssl versions are installed. Run the following command to install the RPM packages:

rpm -ivh --force < filename >

where, filename is the openssl or libopenssl rpm package file.

For example:

rpm -ivh --force openssl-0.9.8h-30.22.21.1.x86_64.rpm rpm -ivh --force libopenssl0_9_8-0.9.8h-30.22.21.1.x86_64.rpm

Why are the remote RACADM and web-based services unavailable after a property change?

It may take a while for the remote RACADM services and the Web-based interface to become available after the iDRAC web server resets.

The iDRAC Web server is reset when:

- The network configuration or network security properties are changed using the iDRAC web user interface.
- The iDRAC.Webserver.HttpsPort property is changed, including when a racadm set -f <config file> changes it.
- The racresetcfg command is used.
- iDRAC is reset.
- A new SSL server certificate is uploaded.

Why is an error message displayed if you try to delete a partition after creating it using local RACADM?

This occurs because the create partition operation is in-progress. However, the partition is deleted after sometime and a message that the partition is deleted is displayed. If not, wait until the create partition operation is completed and then delete the partition.

Permanently setting the default password to calvin

If your system shipped with a unique default iDRAC password but you want to set *calvin* as the default password, you must use the jumpers available on the system board.

CAUTION: Changing the jumper settings permanently changes the default password to *calvin.* You cannot revert to the unique password even if you reset iDRAC to factory settings.

For information about the jumper location and the procedure, see the documentation for your server at https://www.dell.com/ support.

Fls. <u>742</u> Mov. 32

Miscellaneous

Upgrade fails when upgrading to the latest version.

(i) NOTE: 3.30.30.30 is the minimum iDRAC version required to upgrade to 4.00.00.00 / 4.10.10.10 of later build .

After an iDRAC reset, iDRAC GUI may not display all the values.

NOTE: If you reset the iDRAC for some reason, ensure that you wait for at least two minutes after resetting iDRAC to access or modify any settings in iDRAC.

When an OS is installed, hostname may or may not appear/change automatically.

There are two scenarios:

- Scenario 1: iDRAC is not showing the latest hostname once you install an OS. You need to install OMSA or iSM along with the iDRAC to get the hostname reflected.
- Scenario 2: iDRAC had a hostname for a specific OS and another different OS has been installed and still the hostname is appearing as the old hostname without overwriting the hostname. The reason behind, hostname is an information which is coming from the OS, iDRAC only saves the information. If there is a new OS has been installed, iDRAC does not reset the value of the hostname. However, newer versions of the OSs are capable to update the hostname in iDRAC during the 1st OS startup.

How to find an iDRAC IP address for a blade server?

(i) NOTE: The Chassis Management Controller (CMC) option is applicable only for Blade servers.

• Using CMC web interface:

Go to **Chassis** > **Servers** > **Setup** > **Deploy**. In the table that is displayed, view the IP address for the server.

• Using the Virtual Console: Reboot the server to view the iDRAC IP address during POST. Select the "Dell CMC" console in the OSCAR interface to log in to CMC through a local serial connection. CMC RACADM commands can be sent from this connection.

For more information on CMC RACADM commands, see the *Chassis Management Controller RACADM CLI Guide* available at https://www.dell.com/cmcmanuals.

For more information on iDRAC RACADM commands, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/ idracmanuals.

• Using local RACADM

Use the command: racadm getsysinfo For example:

```
$ racadm getniccfg -m server-1
DHCP Enabled = 1
IP Address = 192.168.0.1
Subnet Mask = 255.255.255.0
Gateway = 192.168.0.1
```

• Using LCD:

On the main menu, highlight the server, press the check button, select the required server, and press the check button.



How to find an iDRAC IP address for a blade server ?

i NOTE: The OME-Modular web interface option is applicable only for MX platforms.

• Using OME-Modular web interface:

Go to **Devices** > **Compute**. Select the computer sled and iDRAC IP is displayed as **Management IP**.

- Using OMM Application: see the Dell EMC OpenManage Mobile User's Guide available at https://www.dell.com/ openmanagemanuals
- Using Serial Connection
- Using LCD: On the main menu, highlight the server, press the check button, select the required server, and press the check button.

How to find the CMC IP address related to the blade server?

(i) NOTE: Not applicable for MX platforms.

• From iDRAC web interface:

Go to **iDRAC Settings** > **CMC**. The **CMC Summary** page displays the CMC IP address.

• From the Virtual Console:

Select the "Dell CMC" console in the OSCAR interface to log in to CMC through a local serial connection. CMC RACADM commands can be issued from this connection.

```
$ racadm getniccfg -m chassis
NIC Enabled
                    = 1
                   = 1
DHCP Enabled
                   = 192.168.0.120
Static IP Address
Static Subnet Mask = 255.255.255.0
                    = 192.168.0.1
Static Gateway
Current IP Address = 10.35.155.151
Current Subnet Mask = 255.255.255.0
Current Gateway
                   = 10.35.155.1
Speed
                    = Autonegotiate
Duplex
                    = Autonegotiate
```

(i) NOTE: You can also perform this using remote RACADM.

For more information on CMC RACADM commands, see the *Chassis Management Controller RACADM CLI Guide* available at https://www.dell.com/cmcmanuals.

For more information on iDRAC RACADM commands, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/ idracmanuals.

How to find the OME Modular IP address?

(i) NOTE: Applicable only for MX platforms.

• From iDRAC web interface:

Go to **iDRAC Settings** > Management Module. The Management Module page displays the OME Modular IP address.

How to find iDRAC IP address for rack and tower server?

- From Local RACADM:
 - Use the command racadm getsysinfo.
- From LCD:



On the physical server, use the LCD panel navigation buttons to view the iDRAC IP address. Go to **Setup View** > **View** > **iDRAC IP** > **IPv4** or **IPv6** > **IP**.

From OpenManage Server Administrator:

In the Server Administrator web interface, go to Modular Enclosure > System/Server Module > Main System Chassis/ Main System > Remote Access.

iDRAC network connection is not working.

For blade servers:

- Ensure that the LAN cable is connected to CMC. (not for MX platforms)
- Ensure that NIC settings, IPv4 or IPv6 settings, and either Static or DHCP is enabled for your network.

For rack and tower servers:

- In shared mode, ensure that the LAN cable is connected to the NIC port where the wrench symbol is present.
- In Dedicated mode, ensure that the LAN cable is connected to the iDRAC LAN port.
- Ensure that NIC settings, IPv4 and IPv6 settings and either Static or DHCP is enabled for your network.

iDRAC not accessible in shared LOM

iDRAC may be inaccessible if there are fatal errors in host OS such as BSOD error in Windows. To access iDRAC, reboot the host to recover the connection.

Shared LOM not functional after enabling Link Aggregation Control Protocol (LACP).

The host OS driver for the network adapter must be loaded before LACP is enabled. However, if a passive LACP configuration is in use, the shared LOM may be functional before the host OS driver is loaded. See the switch documentation for LACP configuration.

(i) NOTE: Shared LOM IP of iDRAC is not accessible in pre-boot state when the switch is configured with LACP.

Inserted the blade server into the chassis and pressed the power switch, but it did not power on.

- iDRAC requires up to two minutes to initialize before the server can power on.
- Check CMC, andOME Modular (only for MX platforms) power budget. The chassis power budget may have exceeded.

How to retrieve an iDRAC administrative user name and password?

You must restore iDRAC to its default settings. For more information, see Resetting iDRAC to factory default settings.

How to change the name of the slot for the system in a chassis?

(i) NOTE: Not applicable for MX platforms.

- 1. Log in to CMC web interface and go to Chassis > Servers > Setup.
- 2. Enter the new name for the slot in the row for your server and click Apply.

iDRAC on blade server is not responding during boot.

Remove and reinsert the server.

358 Frequently asked questions



Check CMC (not for MX platforms), and OME Modular (Applicable for MX platforms) web interface to see if iDRAC is displayed as an upgradable component. If it does, follow the instructions in Updating firmware using CMC web interface update the firmware.

(i) NOTE: Update feature not applicable for MX platforms.

If the problem persists, contact technical support.

When attempting to boot the managed server, the power indicator is green, but there is no POST or no video.

This happens due to any of the following conditions:

- Memory is not installed or is inaccessible.
- CPU is not installed or is inaccessible
- Video riser card is missing or not connected properly.

Also, see error messages in iDRAC log using iDRAC web interface or from the server LCD.

Unable to login to iDRAC web interface using Firefox browser on Linux or Ubuntu. Unable to enter the password.

To resolve this issue, reinstall or upgrade the Firefox browser.

Unable to access iDRAC through USB NIC in SLES and Ubuntu

(i) NOTE: In SLES, set the iDRAC interface to DHCP.

In Ubuntu, use the Netplan utility to configure iDRAC interface into DHCP mode. To configure the DHCP:

- 1. Use /etc/netplan/01-netcfg.yaml.
- **2.** Specify Yes for iDRAC DHCP.
- **3.** Apply the configuration.





Figure 5. Configuring iDRAC interface to DHCP mode in Ubuntu

Model, Manufacturer and other properties are not listing for Embedded Network Adapters in Redfish

FRU details for embedded devices will not be displayed. There will not be any FRU object for devices which are embedded on Motherboard. Hence dependent property will not be there.



Use case scenarios

This section helps you in navigating to specific sections in the guide to perform typical use case scenarios.

Topics:

- Troubleshooting an inaccessible managed system
- Obtaining system information and assess system health
- Setting up alerts and configuring email alerts
- Viewing and exporting System Event Log and Lifecycle Log
- Interfaces to update iDRAC firmware
- Performing graceful shutdown
- Creating new administrator user account
- Launching servers remote console and mounting a USB drive
- Installing bare metal OS using attached virtual media and remote file share
- Managing rack density
- Installing new electronic license
- Applying IO Identity configuration settings for multiple network cards in single host system reboot

Troubleshooting an inaccessible managed system

After receiving alerts from OpenManage Essentials, Dell Management Console, or a local trap collector, five servers in a data center are not accessible with issues such as hanging operating system or server. Need to identify the cause to troubleshoot and bring up the server using iDRAC.

Before troubleshooting the inaccessible system, make sure that the following prerequisites are met:

- Enable last crash screen
- Alerts are enabled on iDRAC

To identify the cause, check the following in the iDRAC web interface and re-establish the connection to the system:

NOTE: If you cannot access the iDRAC web interface, go to the server, access the LCD panel, write down the IP address or the host name, and then perform the following operations using iDRAC web interface from your management station:

- Server's LED status Blinking amber or Solid amber.
- Front Panel LCD status or error message Amber LCD or error message.
- Operating system image is seen in the Virtual Console. If you can see the image, reset the system (warm boot) and log in again. If you are able to log in, the issue is fixed.
- Last crash screen.
- Boot capture video.
- Crash capture video.
- Server Health status Red *x* icons for the system components with issues.
- Storage array status Possible array offline or failed
- Lifecycle log for critical events related to system hardware and firmware and the log entries that were logged at the time of system crash.
- Generate Tech Support report and view the collected data.
- Use the monitoring features provided by iDRAC Service Module

Obtaining system information and assess system health

To obtain system information and assess system health:



- In iDRAC Web interface, go to Overview > Summary to view the system information and access various links on this page to asses system health. For example, you can check the health of the chassis fan.
- You can also configure the chassis locator LED and based on the color, assess the system health.
- If iDRAC Service Module is installed, the operating system host information is displayed.

Setting up alerts and configuring email alerts

To set up alerts and configure email alerts:

- 1. Enable alerts.
- 2. Configure the email alert and check the ports.
- 3. Perform a reboot, power off, or power cycle the managed system.
- 4. Send test alert.

Viewing and exporting System Event Log and Lifecycle Log

To view and export lifecycle log and system event log (SEL):

1. In iDRAC Web interface, go to Maintenance > System Event Logs to view SEL and Lifecycle Log to view lifecycle log.

(i) NOTE: The SEL is also recorded in the lifecycle log. Using the filtering options to view the SEL.

- 2. Export the SEL or lifecycle log in the XML format to an external location (management station, USB, network share, and so on). Alternatively, you can enable remote system logging, so that all the logs written to the lifecycle log are also simultaneously written to the configured remote server(s).
- **3.** If you are using the iDRAC Service Module, export the Lifecycle log to OS log.

Interfaces to update iDRAC firmware

Use the following interfaces to update the iDRAC firmware:

- iDRAC Web interface
- Redfish API
- RACADM CLI (iDRAC_) and CMC (not applicable for MX platforms))
- Dell Update Package (DUP)
- CMC (not applicable for MX platforms)OME Modular (applicable only for MX platforms) Web interface
- Lifecycle Controller-Remote Services
- Lifecycle Controller
- Dell Remote Access Configuration Tool (DRACT)

Performing graceful shutdown

To perform graceful shutdown, in iDRAC Web interface, go to one of the following locations:

At Dashboard select Graceful Shutdown and click Apply.

For more information, see the *iDRAC Online Help*.

Creating new administrator user account

You can modify the default local administrator user account or create a new administrator user account. To modify the local administrator user account, see Modifying local administrator account settings.

To create a new administrator account, see the following sections:

• Configuring local users



- Configuring active directory users
- Configuring generic LDAP users

Launching servers remote console and mounting a USB drive

To launch the remote console and mount a USB drive:

- 1. Connect a USB flash drive (with the required image) to the management station.
- 2. Use the following method to launch virtual console through the iDRAC Web Interface:
- Go to **Dashboard > Virtual Console** and click **Launch Virtual Console**.

The Virtual Console Viewer is displayed.

- 3. From the File menu, click Virtual Media > Launch Virtual Media.
- Click Add Image and select the image that is located on the USB flash drive. The image is added to the list of available drives.
- 5. Select the drive to map it. The image on the USB flash drive is mapped to the managed system.

Installing bare metal OS using attached virtual media and remote file share

See the Deploying operating system using remote file share section.

Managing rack density

Before you install additional servers in a rack, you must determine the remaining capacity in the rack.

To assess the capacity of a rack to add additional servers:

- 1. View the current power consumption data and historical power consumption data for the servers.
- 2. Based on the data, power infrastructure and cooling system limitations, enable the power cap policy and set the power cap values.
 - **NOTE:** It is recommended that you set a cap close to the peak, and then use that capped level to determine how much capacity is remaining in the rack for adding more servers.

Installing new electronic license

See License operations for more information.

Applying IO Identity configuration settings for multiple network cards in single host system reboot

If you have multiple network cards in a server that is part of a Storage Area Network (SAN) environment and you want to apply different virtual addresses, initiator and target configuration settings to those cards, use the I/O Identity Optimization feature to reduce the time in configuring the settings. To do this:

- 1. Make sure that BIOS, iDRAC, and the network cards are updated to the latest firmware version.
- 2. Enable IO Identity Optimization.
- 3. Export the Server Configuration Profile (SCP) file from iDRAC.
- 4. Edit the I/O Identity optimization settings in the SCP file.

5. Import the SCP file to iDRAC.





Dell PowerEdge RAID Controller 11 User's Guide

PERC H755, H750, H355, and H350 Controller Series

Regulatory Model: UCPA-1101, UCPF-1100, UCPF-1110, UCPN-1100, UCSA-1111, UCSF-1100, and UCSM-1100 September 2023 Rev. A05





Notes, cautions, and warnings

(i) NOTE: A NOTE indicates important information that helps you make better use of your product.

CAUTION: A CAUTION indicates either potential damage to hardware or loss of data and tells you how to avoid the problem.

WARNING: A WARNING indicates a potential for property damage, personal injury, or death.

© 2020-2022 Dell Inc. or its subsidiaries. All rights reserved. Dell Technologies, Dell, and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be trademarks of their respective owners.



Chanter 1: Dell Technologies PowerEdge RAID Controller 11	8
Eastures of PERC H755 adapter	9 9
Features of PERC H755 front SAS	9
Features of PERC H755N front NVMe	
Features of PERC H755 MX adapter	10
Features of PERC H750 adapter SAS	11
Features of PERC H355 adapter SAS	11
Features of PERC H355 front SAS	12
Features of PERC H350 adapter SAS	
Features of PERC H350 Mini Monolithic SAS	13
Operating systems supported by PERC 11 cards	13
Technical specifications of PERC 11 cards	10 14
Thermal specifications	
Chanter 2: Applications and User Interfaces supported by PERC 11	18
Comprehensive Embedded Management	18
Dell OpenManage Storage Management	18
Human Interface Infrastructure Configuration Litility	18
The PERC Command Line Interface	19
Chapter 3: Eastures of PowerEdge PAID Controller 11	20
Controller features	20
Non-Volatile Memory Express	20
Onal Security Management	
Hardware Root of Trust	
1 MR I/O	
Disk roaming	21
EastPath	22 20
Non-RAID disks	
Physical disk nower management	23
Profile Management	23
Secure firmware undate	23
Spandump	20
Virtual disk features	
Virtual disk write cache policy	24
Virtual disk write eache policy	27 24
Virtual disk read cache policy	
Virtual disk initialization	25
Full initialization	20
	25
Reconfigure virtual disks	20 26
Reckaround operations	
	70
Background initialization	

Contents 3



Consistency checks	
Hard drive features	
Self-Encrypting Disks	
Instant secure erase	
4 KB sector disk drives	
Fault tolerance	
The SMART feature	29
Patrol Read	
Physical disk failure detection	
Controller cache	
Battery Transparent Learn Cycle	
Linux operating system device enumeration	32
Chapter 4: Install and remove a PERC 11 card	
Safety instructions	34
Before working inside your system	
Remove the PERC H755 adapter	35
Install the PERC H755 adapter	
Remove the PERC H755 front SAS card	
Install the PERC H755 front SAS card	38
Remove the PERC H755N front NVMe card	
Install the PERC H755N front NVMe card	
Remove the PERC H755 MX adapter	42
Install the PERC H755 MX adapter	
Remove the PERC H750 adapter SAS	45
Install the PERC H750 adapter SAS	
Remove the PERC H355 adapter SAS	
Install the PERC H355 adapter SAS	
Remove the PERC H355 front SAS	
Install the PERC H355 front SAS card	
Remove the PERC H350 adapter SAS	
Install the PERC H350 adapter SAS	
Remove PERC H350 Mini Monolithic SAS	
Install PERC H350 Mini Monolithic SAS	55
Chapter 5: Driver support for PERC 11	57
Creating the device driver media	
Download and save PERC 11 drivers from the support site	
Download and save PERC 11 drivers from the Dell Systems Service and Diagnostic Tools	
Windows driver installation	
Install PERC 11 driver while newly installing the Windows Server 2016 and later	
Install PERC 11 driver on which the windows Server 2016 is already installed and later	
Update PERC 11 driver that runs on windows Server 2016 or later	
Initial or update a PPM driver package using the KMOD support	
Install or update a REIVI univer package using the KMP support	
Instantion update a REIVI univer package using the RIVIE Support	0U 61
Load the driver write thatalling an operating system	
Chapter 6: Firmware	62



Upgrade firmware controller using Dell Update Package (DUP).....

Chapter 7: Manage PERC 11 controllers using HII configuration utility	63
Enter the PERC 11 HII configuration utility	
Exit the PERC 11 HII configuration utility	
Navigate to Dell PERC 11 configuration utility	64
View the HII Configuration utility dashboard	64
Configuration management	65
Auto Configure RAID 0	65
Create virtual disks	65
Create profile based virtual disk	
View disk group properties	67
Convert to Non-RAID disk	67
Delete configurations	67
Controller management	
Clear controller events	68
Save controller events	
Save debug log	68
Enable security	68
Disable security	68
Change security settings	69
Restore factory default settings	69
Auto configure behavior	
Manage controller profile	
Advanced controller properties	70
Virtual disk management	73
Virtual disk numbering	73
Configure Virtual Disks	75
Perform expand virtual disk operation	75
Perform consistency check	76
Physical disk management	
View physical disk properties	76
Cryptographic erase	77
Physical disk erase	
Assigning a global hot spare	78
Assigning a dedicated hot spare	
Convert to Non–RAID disk	79
Hardware components	80
View battery properties	80
View physical disks associated with an enclosure	80
Security key management in HII configuration utility	81
Chapter 8: Security key and RAID management	
Security key implementation	
Local Key Management	
Create a security key	
Change Security Settings	
Disable security key	
Create a secured virtual disk	

Contents 5



Secure a non-RAID disk	
Secure a pre-existing virtual disk	
Import a secured non-RAID disk	
Import a secured virtual disk	
Dell Technologies OpenManage Secure Enterprise Key Manager	
Supported controllers for OpenManage Secure Enterprise Key Manager	85
Manage enterprise key manager mode	
Disable enterprise key manager mode	86
Manage virtual disks in enterprise key manager mode	
Manage non–RAID disks in enterprise key manager mode	
Transition of drives from local key management to enterprise key management (without supported firmware for PERC and iDRAC)	86
Migrate of drives from local key management to enterprise key management (with supported firmware for PERC and iDRAC)	ا 87
Chapter 9: Troubleshooting issues in PERC11 cards	88
Single virtual disk performance or latency in hypervisor configurations	
Configured disks removed or not accessible error message	
Dirty cache data error message	89
Discovery error message	
Drive Configuration Changes Error Message	
Windows operating system installation errors	
Firmware fault state error message	
Foreign configuration found error message	
Foreign configuration not found in HII	
Degraded state of virtual disks	90
Memory errors	
Preserved Cache State	
Security key errors	91
Secured foreign import errors	91
Failure to select or configure non Self-Encrypting Disks non-SED	91
Failure to delete security key	92
Failure of Cryptographic Erase on encryption-capable physical disks	
General issues	
PERC card has yellow bang in Windows operating system device manager	92
PERC card not seen in operating systems	92
Issues in controller, battery, and disk when operating at low temperature	
Physical disk issues	92
Physical disk in failed state	
Unable to rebuild a fault tolerant virtual disk	93
Fatal error or data corruption reported	
Multiple disks are inaccessible	93
Rebuilding data for a failed physical disk	93
Virtual disk fails during rebuild using a global hot spare	
Dedicated hot spare disk fails during rebuild	94
Redundant virtual disk fails during reconstruction	94
Virtual disk fails rebuild using a dedicated hot spare	94
Physical disk takes a long time to rebuild	94
Drive removal and insertion in the same slot generates a foreign configuration event	94
SMART errors	



Smart error detected on a non-RAID disk	
Smart error detected on a physical disk in a non-redundant virtual disk	
Smart error detected on a physical disk in a redundant virtual disk	
Replace member errors	
Source disk fails during replace member operation	
Target disk fails during replace member operation	
A member disk failure is reported in the virtual disk which undergoes replace mem	nber operation96
Linux operating system errors	
Virtual disk policy is assumed as write-through	
Unable to register SCSI device error message	
Drive indicator codes	
HII error messages	
Unhealthy Status of the drivers	
Rebuilding a drive during full initialization	
System reports more drive slots than what is available	
World Wide Number on drive sticker is not the same in applications	
Backplane firmware revision not changing in PERC interfaces after an update	
Chapter 10: Appendix RAID description	
Summary of RAID levels	
RAID 10 configuration	
RAID terminology	
Disk striping	
Disk mirroring	
Spanned RAID levels	
Spanned RAID levels Parity data	
Spanned RAID levels Parity data Chapter 11: Getting help	104
Spanned RAID levels Parity data Chapter 11: Getting help Recycling or End-of-Life service information	104
Spanned RAID levels Parity data Chapter 11: Getting help Recycling or End-of-Life service information Contacting Dell	104
Spanned RAID levels Parity data Chapter 11: Getting help Recycling or End-of-Life service information Contacting Dell Locating the Express Service Code and Service Tag	104
Spanned RAID levels Parity data Chapter 11: Getting help Recycling or End-of-Life service information Contacting Dell Locating the Express Service Code and Service Tag Receiving automated support with SupportAssist	
Spanned RAID levels Parity data Chapter 11: Getting help Recycling or End-of-Life service information Contacting Dell Locating the Express Service Code and Service Tag Receiving automated support with SupportAssist Chapter 12: Documentation resources.	104



Dell Technologies PowerEdge RAID Controller 11

Dell Technologies PowerEdge RAID Controller 11, or PERC 11 is a series of RAID disk array controllers made by Dell for its PowerEdge servers. The PERC 11 series consists of the PERC H755 adapter, PERC H755 front SAS, PERC H755N front NVMe, PERC H750 adapter SAS, PERC H755 MX adapter, PERC H355 adapter SAS, PERC H355 front SAS, PERC H350 adapter SAS, and PERC H350 Mini Monolithic SAS cards that have the following characteristics:

- Provides reliability, high performance, and fault-tolerant disk subsystem management
- Offers RAID control capabilities including support for RAID levels 0, 1, 5, 6, 10, 50, 60
- Complies with Serial Attached SCSI (SAS) 3.0 providing up to 12 Gb/sec throughput
- Supports Dell-qualified Serial Attached SCSI (SAS), SATA hard drives, solid state drives (SSDs), and PCIe SSD (NVMe)
- Supports drive speeds of 8 GT/s and 16 GT/s at maximum x2 lane width for NVMe drives.
- **NOTE:** Mixing disks of different speeds (7,200 RPM, 10,000 RPM, or 15,000 RPM) and bandwidth (3 Gbps, 6 Gbps, or 12 Gbps) while maintaining the same drive type (SAS or SATA) and technology (hard drive or SSD) is supported.
- **NOTE:** Mixing NVMe drives with SAS and SATA is not supported. Also, mixing hard drives and SSDs in a virtual disk is not supported.
- **NOTE:** PERC H750 adapter SAS, PERC H355 front SAS, PERC H355 adapter SAS, PERC H350 adapter SAS, and PERC H350 Mini Monolithic SAS do not support NVMe drives.
- **NOTE:** RAID levels 5, 6, 50, and 60 are not supported on PERC H355 adapter SAS, PERC H355 front SAS, PERC H350 adapter SAS, and PERC H350 Mini Monolithic SAS.
- **NOTE:** PERC H350 Mini Monolithic SAS has form factor variations (Low Profile) for specific platforms. For more information, see your platform manuals.
- () NOTE: For the safety, regulatory, and ergonomic information that is associated with these devices, and for more information about the Integrated Dell Remote Access Controller (iDRAC) or Lifecycle Controller (LC) remote management, see your platform documentation.

Topics:

- Features of PERC H755 adapter
- Features of PERC H755 front SAS
- Features of PERC H755N front NVMe
- Features of PERC H755 MX adapter
- Features of PERC H750 adapter SAS
- Features of PERC H355 adapter SAS
- Features of PERC H355 front SAS
- Features of PERC H350 adapter SAS
- Features of PERC H350 Mini Monolithic SAS
- Operating systems supported by PERC 11 cards
- Technical specifications of PERC 11 cards
- Thermal specifications



Features of PERC H755 adapter

This section describes the features of PERC H755 adapter.



Figure 1. Features of PERC H755 adapter

- 1. Heatsink
- 3. Battery
- 5. Backplane connector B

- 2. PCle connector
- 4. Backplane connector A
- 6. Battery cable connector

Features of PERC H755 front SAS

This section describes the features of PERC H755 front SAS.



Figure 2. Features of PERC H755 front SAS

1. Battery

2. Backplane connector A

Dell Technologies PowerEdge RAID Controller 11 9

FIS. 760 Mov. 32

- 3. Power card edge connector
- 5. PCle input connector
- 7. Battery cable connector

- 4. Heatsink
- 6. Backplane connector B

Features of PERC H755N front NVMe



Figure 3. Features of PERC H755N front NVMe

- 1. Battery
- 3. Power card edge connector
- 5. Backplane connector A
- 7. Battery cable connector

6. Backplane connector B

4. Heatsink

2. PCIe cable connector

Features of PERC H755 MX adapter



Figure 4. Features of PERC H755 MX adapter

- 1. Battery under cover
- 3. PCIe cable connector
- 5. Backplane connector B

- 2. Heatsink
- 4. Backplane connector A

10 Dell Technologies PowerEdge RAID Controller 11



Features of PERC H750 adapter SAS



Figure 5. Features of PERC H750 adapter SAS

- 1. Heat sink
- 3. Battery cable connector

Battery
 Backplane connector A

5. PCle connector

Features of PERC H355 adapter SAS



Figure 6. Features of PERC H355 adapter SAS

- 1. Heat sink
- 3. Backplane connector A

Backplane connector B
 PCle connector

Dell Technologies PowerEdge RAID Controller 11 11



Features of PERC H355 front SAS



Figure 7. Features of H355 front SAS

- 1. PCle input connector
- 3. Backplane connector B
- 5. Power card edge connector

- 2. Heat sink
- 4. Backplane connector A

Features of PERC H350 adapter SAS



Figure 8. PERC H350 adapter SAS

- 1. Heat sink
- 2. Backplane connector A
- 3. PCle connector



Features of PERC H350 Mini Monolithic SAS



Figure 9. PERC H350 Mini Monolithic SAS

- 1. SAS cable connection
- 2. Heat sink

Operating systems supported by PERC 11 cards

See Dell Technologies Enterprise operating systems support for a list of supported operating systems by a specific server for the PERC 11 cards.

() NOTE: For the latest list of supported operating systems and driver installation instructions, see the operating system documentation at Operating Systems Documentation. For specific operating system service pack requirements, see the Drivers and Downloads section on the support site.



Technical specifications of PERC 11 cards

The following table lists the specifications of PERC 11 cards:

Table 1. Technical specifications of PERC 11 cards

Feature	PERC H755 adapter	PERC H755 front SAS	PERC H755N front NVMe	PERC H755 MX adapter	PERC H750 adapter SAS
RAID levels	0, 1, 5, 6, 10, 50, 60	0, 1, 5, 6, 10, 50, 60	0, 1, 5, 6, 10, 50, 60	0, 1, 5, 6, 10, 50 ,60	0, 1, 5, 6, 10, 50 ,60
Non-RAID	Yes	Yes	Yes	Yes	Yes
Enclosures per port	Not applicable				
Processor	Broadcom RAID- on-chip, SAS3916 chipset	Broadcom RAID- on-chip, SAS3916 chipset	Broadcom RAID-on- chip, SAS3916 chipset	Broadcom RAID- on-chip, SAS3916 chipset	Broadcom RAID-on- chip, SAS3916 chipset
Battery backup unit	Yes	Yes	Yes	Yes	Yes
Local Key Management security	Yes	Yes	Yes	Yes	Yes
Controller queue depth	5120	5120	5120	5120	5120
Secure enterprise key manager security	Yes	Yes	Yes	No	Yes
Non-volatile cache	Yes	Yes	Yes	Yes	Yes
Cache memory	8 GB DDR4 2666 MT/s cache	8 GB DDR4 2666 MT/s cache	8 GB DDR4 2666 MT/s cache	8 GB DDR4 2666 MT/s cache	8 GB DDR4 2666 MT/s cache
Cache function	Write back, write through, no read ahead, and read ahead	Write back, write through, no read ahead, and read ahead	Write back, write through, no read ahead, and read ahead	Write back, write through, no read ahead, and read ahead	Write back, write through, no read ahead, and read ahead
Max no of VDs in RAID mode	240	240	240	240	240
Max no of disk groups	240	240	240	240	240
Max no of VDs per disk group	16	16	16	16	16
Hot swap devices supported	Yes	Yes	Yes	Yes	Yes
Autoconfig	Yes	Yes	Yes	Yes	Yes
Hardware XOR engine	Yes	Yes	Yes	Yes	Yes
Online capacity expansion	Yes	Yes	Yes	Yes	Yes
Dedicated and global hot spare	Yes	Yes	Yes	Yes	Yes
Drives types	3 Gbps SATA, 6 Gbps SATA/SAS, and 12 Gbps SAS,	3 Gbps SATA, 6 Gbps SATA/SAS, and 12 Gbps SAS	Gen3 (8 GT/s) and Gen4 (16 GT/s) NVMe	3 Gbps SATA, 6 Gbps SATA/ SAS, and 12	3 Gbps SATA, 6 Gbps SATA/SAS, and 12 Gbps SAS

14 Dell Technologies PowerEdge RAID Controller 11



Feature	PERC H755 adapter	PERC H755 front SAS	PERC H755N front NVMe	PERC H755 MX adapter	PERC H750 adapter SAS
	Gen3 (8 GT/s), and Gen4 (16 GT/s) NVMe			Gbps SAS, Gen3 (8 GT/s), and Gen4 (16 GT/s) NVMe	
VD strip size	64 KB, 128 KB, 256 KB, 512 KB, 1 MB	64 KB, 128 KB, 256 KB, 512 KB, 1 MB	64 KB, 128 KB, 256 KB, 512 KB, 1 MB	64 KB, 128 KB, 256 KB, 512 KB, and 1 MB	64 KB, 128 KB, 256 KB, 512 KB, 1 MB
PCle support	Gen 4	Gen 4	Gen 4	Gen 4	Gen 4
SAS/SATA maximum drive support	 Without SAS Expander: 16 drives per controller With SAS Expander: Limited by platform offerings 	 Without SAS Expander: 16 drives per controller With SAS Expander: Limited by platform offerings 	Not applicable	Limited by platform: 8 drives per controller	 Without SAS Expander: 8 drives per controller With SAS Expander: Limited by platform offerings
NVMe maximum drive support	 Without PCle Switch Expander: 8 drives per controller With PCle Switch Expander: Limited by platform offerings 	Not applicable	 Without PCle Switch Expander: 8 drives per controller With PCle Switch Expander: Limited by platform offerings 	Limited by platform:8 drives per controller	Not applicable

Table 1. Technical specifications of PERC 11 cards (continued)

NOTE: PERC H755 adapter and PERC H755 MX supports either SAS, SATA, or NVMe drives depending on the backplane/ server configuration.

NOTE: PERC controller supports only conventional magnetic recording (CMR) drives, and does not support shingled magnetic recording (SMR) drives.

(i) NOTE: PERC H755 family of controllers currently support SEKM starting with firmware version 52.14.0-3901.

(i) NOTE: For information about the number of drives in a disk group per virtual disk, see Summary of RAID levels

(i) NOTE: As 14G PowerEdge Servers do not support Gen 4 speeds, PERC H750 adapter SAS will downtrain to Gen 3 speeds.

Table 2. Technical specifications of PERC 11 cards

Feature	PERC H355 adapter SAS	PERC H355 front SAS	PERC H350 adapter SAS	PERC H350 Mini Monolithic SAS
RAID levels	0, 1, 10	0, 1, 10	0, 1, 10	0, 1, 10
Non-RAID	Yes	Yes	Yes	Yes
Enclosures per port	Not applicable	Not applicable	Not applicable	Not applicable
Processor	Broadcom RAID- onchip, SAS3816 chipset	Broadcom RAID- onchip, SAS3816 chipset	Broadcom RAID- onchip, SAS3816 chipset	Broadcom RAID-onchip, SAS3816 chipset
Battery backup unit	No	No	No	No



Table 2. Technical specifications of PERC 11 cards (continued)

Feature	PERC H355 adapter SAS	PERC H355 front SAS	PERC H350 adapter SAS	PERC H350 Mini Monolithic SAS
Local Key Management security	No	No	No	No
Controller queue depth	1536	1536	1536	1536
Secure enterprise key manager security	No	No	No	No
Non-volatile cache	No	No	No	No
Cache memory	Not applicable	Not applicable	Not applicable	Not applicable
Cache function	Write through, no read ahead	Write through, no read ahead	write through, no read ahead	write through, no read ahead
Max no of VDs in RAID mode	32	32	32	32
Max no of disk groups	32	32	32	32
Max no of VDs per disk group	16	16	16	16
Hot swap devices supported	Yes	Yes	Yes	Yes
Autoconfig	Yes	Yes	Yes	Yes
Hardware XOR engine	Yes	Yes	Yes	Yes
Online capacity expansion	Yes	Yes	Yes	Yes
Dedicated and global hot spare	Yes	Yes	Yes	Yes
Drives types	3 Gbps SATA, 6 Gbps SATA/SAS, and 12 Gbps SAS, Gen3 (8 GT/s)	3 Gbps SATA, 6 Gbps SATA/SAS, and 12 Gbps SAS, Gen3 (8 GT/s)	3 Gbps SATA, 6 Gbps SATA/SAS, and 12 Gbps SAS, Gen3 (8 GT/s)	3 Gbps SATA, 6 Gbps SATA/SAS, and 12 Gbps SAS, Gen3 (8 GT/s)
VD strip size	64 KB	64 KB	64 KB	64 KB
PCle support	Gen 4	Gen 4	Gen 4	Gen 4
SAS/SATA maximum drive support	 Without SAS Expander: 16 With SAS Expander: Limited by platform offering 	 Without SAS Expander: 16 With SAS Expander: Limited by platform offering 	 Without SAS Expander: 8 drives per controller With SAS Expander: Limited by platform offering 	 Without SAS Expander: 8 drives per controller With SAS Expander: Limited by platform offering
NVMe maximum drive support	Not applicable	Not applicable	Not applicable	Not applicable

NOTE: As 14G PowerEdge Servers do not support Gen 4 speeds, PERC H350 adapter SAS and PERC H350 Mini Monolithic SAS will down train to Gen 3 speeds.

Thermal specifications

PERC 11 Controllers have an operating temperature range of 0C to 55C. System ambient temperatures may be less than or greater than these values.

16 Dell Technologies PowerEdge RAID Controller 11



() NOTE: PERC Controllers may raise erroneous Battery, Disk, and Controller temperature errors if the controller is operating below the operational temperature range.



Applications and User Interfaces supported by PERC 11

PERC 11 card Management applications include the Comprehensive Embedded Management (CEM), Dell OpenManage Storage Management, The Human Interface Infrastructure (HII) configuration utility, and The PERC Command Line Interface (CLI). They enable you to manage and configure the RAID system, create and manage multiple disk groups, control and monitor multiple RAID systems, and provide online maintenance.

Topics:

- Comprehensive Embedded Management
- Dell OpenManage Storage Management
- Human Interface Infrastructure Configuration Utility
- The PERC Command Line Interface

Comprehensive Embedded Management

Comprehensive Embedded Management (CEM) is a storage management solution for Dell systems that enables you to monitor the RAID and network controllers installed on the system using iDRAC without an operating system installed on the system.

Using CEM enables you to do the following:

- Monitor devices with and without an operating systems installed on the system
- Provide a specific location to access monitored data of the storage devices and network cards
- Allows controller configuration for all PERC 11 cards
- **NOTE:** If you boot the system to HII (F2) or Lifecycle Controller (F10), then you cannot view the PERC cards on the CEM UI. The PERC cards are displayed on the CEM UI only after the system boot is complete.

(i) NOTE: It is not recommended that you create more than 8 VDs simultaneously with CEM.

Dell OpenManage Storage Management

Dell OpenManage Storage Management is a storage management application for Dell systems that provides enhanced features for configuring locally attached RAID disk storage. The Dell OpenManage storage management application enables you to perform controller and enclosure functions for all supported RAID controllers and enclosures from a single graphical or Command Line Interface (CLI). The User Interface (UI) is wizard-driven with features for novice and advanced users, and detailed online help. Using the Dell OpenManage storage management application, you can protect your data by configuring data-redundancy, assigning hot spares, or rebuilding failed physical disks. The fully featured CLI, which is available on select operating systems, allows you to perform RAID management tasks either directly from the console or through scripting.

(i) NOTE: For more information, see the Dell OpenManage Storage Management User's Guide at OpenManage Manuals

Human Interface Infrastructure Configuration Utility

The Human Interface Infrastructure (HII) configuration utility is a storage management application integrated into the system BIOS <F2>. It is used to configure and manage your Dell PowerEdge RAID Controller (PERC) virtual disks, and physical disks. This utility is independent of the operating system.

(i) NOTE: The BIOS configuration utility <Ctrl> <R> is not supported on PERC 11 cards.

18 Applications and User Interfaces supported by PERC 11



The PERC Command Line Interface

The PERC Command Line Interface (CLI) is a storage management application. This utility allows you to set up, configure, and manage your Dell PowerEdge RAID Controller (PERC) by using the Command Line Interface (CLI).

(i) NOTE: For more information, see Dell PowerEdge RAID Controller CLI Reference Guide at Storage Controllers Manuals.

Applications and User Interfaces supported by PERC 11 19



Features of PowerEdge RAID Controller 11

Topics:

- Controller features
- Virtual disk features
- Virtual disk initialization
- Reconfigure virtual disks
- Background operations
- Hard drive features
- Fault tolerance

Controller features

This section lists the following controller features supported on Dell Technologies PowerEdge RAID Controller 11 cards in detail:

- Non-Volatile Memory Express
- Opal Security Management
- Hardware Root of Trust
- 1 MB I/O
- Auto Configure RAID 0
- Disk roaming
- FastPath
- Non–RAID disks
- Physical disk power management
- Profile Management
- Secure firmware update
- Snapdump

Non-Volatile Memory Express

Non-Volatile Memory Express (NVMe) is a standardized, high-performance host controller interface and a storage protocol for communicating with non-volatile memory storage devices over the peripheral component interconnect express (PCle) interface standard. The PERC 11 controller supports up to 8 direct-attach NVMe drives. The PERC 11 controller is a PCle endpoint to the host, a PowerEdge server, and configured as a PCle root complex for downstream PCle NVMe devices connected to the controller.

NOTE: The NVMe drive on the PERC 11 controller shows up as a SCSI disk in the operating system, and the NVMe command line interface will not work for the attached NVMe drives.

Conditions under which a PERC supports an NVMe drive

- In NVMe devices the namespace identifier (NSID) with ID 1, which is (NSID=1) must be present.
- In NVMe devices with multiple namespace(s), you can use the drive capacity of the namespace with NSID=1.
- The namespace with NSID=1 must be formatted without protection information and cannot have the metadata enabled.
- PERC supports 512-bytes or 4 KB sector disk drives for NVMe devices.

20 Features of PowerEdge RAID Controller 11



Drive repair for NVMe initialization failure

If an NVME drive fails to initialize, the drive that is connected to PERC can be corrected in HII. The NVME initialization errors in the drives are listed as correctable and non-correctable errors in HII.

Repair drives with correctable NVMe initialization errors

Repair the drives with correctable NVMe initialization errors in HII to enable the drives to work properly.

About this task

Repairs can lead to permanent data loss in drives. Also, certain types of repairs can take a long time.

Steps

- 1. Log in to HII.
- 2. Go to Main Menu > Hardware Components > Enclosure Management. The drives with correctable and non-correctable errors are listed.
- **3.** Select the drive and click **Repair**. If the repair is successful, the drive is listed under physical drives and removed from the correctable error list. If the drive has other correctable errors, the drive is listed again in the correctable errors list.
- 4. If the repair is not successful, click **Repair** again.

(i) NOTE: In case you want to stop the repair, stop the repair from the Ongoing repairs list.

If the error is still not resolved or if the drive has other non-correctable errors, the drive is moved to the non-correctable error list.

Opal Security Management

Opal Security Management of Opal SED drives requires security key management support. You can use the application software or The Integrated Dell Remote Access Controller (iDRAC) to generate the security key that is set in the Opal drives and used as an authentication key to lock and unlock the Opal drives.

Hardware Root of Trust

Hardware RoT (RoT) builds a chain of trust by authenticating all the firmware components prior to its execution, and it permits only the authenticated firmware to perform and be flashed. The controller boots from an internal boot ROM (IBR) that establishes the initial root of trust and this process authenticates and builds a chain of trust with succeeding software using this root of trust.

1 MB I/O

PERC 11 controllers support a 1 MB I/O feature; if the capacity of I/O frame is greater than 1 MB, the I/O frame is broken into smaller chunks.

Autoconfigure RAID 0

The Autoconfigure RAID 0 feature creates a single drive RAID 0 on each hard drive that is in the ready state. For more information, see Auto Configure RAID 0.

NOTE: The Autoconfigure RAID 0 feature is not supported on PERC H355 adapter SAS, PERC H355 front SAS, PERC H350 adapter SAS, and PERC H350 Mini Monolithic SAS.



Autoconfigure behavior

The autoconfigure behavior automatically configures unconfigured drives during reboot and hot insertion. Unconfigured drives are configured according to the settings; but the configured drives remain unaffected. PERC 11 supports **Off and Non-RAID** settings.

Table 3. Autoconfigure behavior settings

Settings	Description
Off	Autoconfigure behavior is turned off.
Non-RAID	Unconfigured drives are configured as non-RAID disk during boot or during hot insertion; all the configured drives remain unaffected.
Off to Non-RAID disk	Unconfigured drives are converted to non-RAID disks; all the configured drives remain unaffected.
Non-RAID disk to Off	Unconfigured drives remain unconfigured good; all the configured drives remain unaffected.

NOTE: PERC H355 front SAS, PERC H355 adapter SAS, PERC H350 adapter SAS, and PERC H350 Mini Monolithic SAS converts an unconfigured good drive to non-RAID only if the drive has never been used before by that specific PERC.

Disk roaming

Disk roaming is when a physical disk is moved from one cable connection or backplane slot to another on the same controller. The controller automatically recognizes the relocated physical disks and logically places them in the virtual disks that are part of the disk group. If the physical disk is configured as a non-RAID disk, then the relocated physical disk is recognized as a non-RAID disk by the controller.

CAUTION: It is recommended that you perform disk roaming when the system is turned off.

CAUTION: Do not attempt disk roaming during RAID level migration (RLM) or online capacity expansion (OCE). This causes loss of the virtual disk.

Using disk roaming

About this task

Perform the following steps to use disk roaming:

Steps

- 1. Turn off the power to the system, physical disks, enclosures, and system components.
- 2. Disconnect power cables from the system.
- 3. Move the physical disks to desired positions on the backplane or the enclosure.
- 4. Perform a safety check. Make sure the physical disks are inserted properly.
- 5. Turn on the system.

Results

The controller detects the RAID configuration from the configuration data on the physical disks.

FastPath

FastPath is a feature that improves application performance by delivering high I/O per second (IOPs) for solid-state drives (SSDs). The PERC 11 series of cards support FastPath.

To enable FastPath on a virtual disk, the cache policies of the RAID controller must be set to write-through and no read ahead. This enables FastPath to use the proper data path through the controller based on command (read/write), I/O size, and RAID type. For optimal solid-state drive performance, create virtual disks with strip size of 64 KB.

22 Features of PowerEdge RAID Controller 11



Non-RAID disks

A non-RAID disk is a single disk to the host, and not a RAID volume. The only supported cache policy for non-RAID disks is Write-Through.

Physical disk power management

Physical disk power management is a power-saving feature of PERC 11 series cards. The feature allows disks to be spun down based on disk configuration and I/O activity. The feature is supported on all rotating SAS and SATA disks, and includes unconfigured and hot-spare disks. The physical disk power management feature is disabled by default. You can enable the feature in the Dell Open Manage Storage Management application or in the Human Interface Infrastructure (HII) configuration utility. For more information on HII configuration and physical disk power management, see Enabling physical disk power management. For more information on using the Dell Open Manage Storage Management application, see the Dell OpenManage documentation at OpenManage Manuals.

Profile Management

PERC 11 supports the PD240 and PD64 profiles. It defines controller queue depth and the maximum number of physical and virtual disks.

Table 4. Supported profile on PERC 11

Feature	PD240	PD64
Controller	PERC H755 front SAS, PERC H755 MX adapter, and PERC H750 adapter SAS	PERC H355 front SAS, PERC H355 adapter SAS, PERC H350 adapter SAS, and PERC H350 Mini Monolithic SAS
Maximum virtual disk supported	240	32
Controller queue depth	5120	1536

Secure firmware update

This feature provides a cryptographic method of updating the firmware using an RSA encryption-decryption algorithm.

Only Dell-certified firmware is supported on your PERC controller.

Snapdump

The Snapdump feature provides the Dell support team with the debug information which can help to find the cause of firmware failure. In the instance of firmware failures, the firmware collects the logs and information at the time of failure, which are stored in a compressed file called a snapdump.

Snapdumps are also generated manually to provide additional debug information. When a snapdump is generated, it is stored in the controller's cache memory. This means in the event of a power loss the controller will offload the snapdump as part of its cache preservation mechanism. Snapdumps are preserved by default through four reboots before its deleted.

To generate a snapdump, change the snapdump, delete a snapdump, and to download a stored snapdump settings, see Dell PowerEdge RAID Controller CLI Reference Guide at Storage Controllers Manuals.

Virtual disk features

This section lists the following virtual disk features supported on PERC 11 cards in detail:

- Virtual disk read cache policies
- Virtual disk write cache policies
- Virtual disk migration



- Virtual disk initialization
- Reconfiguration of virtual disks
- Background operations

Virtual disk write cache policy

The write cache policy of a virtual disk determines how the controller handles writes to the virtual disk.

Table 5. Write cache policies

Feature	Description
Write-back	The controller sends a data transfer completion signal to the host when the controller cache has received all the data in a transaction. The controller then writes the cached data to the storage device in the background.
Write-through	The controller sends a data transfer completion signal to the host system when the disk subsystem has received all the data in a transaction.

() NOTE: All RAID volumes are presented as write-through to the operating system (Windows and Linux) independent of the actual write cache policy of the virtual disk. PERC cards manage the data in cache independently of the operating system or any applications.

NOTE: Use the Dell OpenManage storage management application or the HII Configuration Utility to view and manage virtual disk cache settings.

Conditions under which write-back is employed

Write-back caching is used under all conditions in which the battery is present and in good condition.

Conditions under which forced write-back with no battery is employed

CAUTION: It is recommended that you use a power backup system when forcing write-back to ensure there is no loss of data if the system suddenly loses power.

Write-back mode is available when you select force write-back with no battery. When forced write-back mode is selected, the virtual disk is in write-back mode even if the battery is not present.

Virtual disk read cache policy

The read policy of a virtual disk determines how the controller handles reads to that virtual disk.

Table 6. Read policies

Feature	Description
Read ahead	Allows the controller to read sequentially ahead of requested data and to store the additional data in cache memory, anticipating that the data is required soon. This speeds up reads for sequential data, but there is slight improvement when accessing random data.
No read ahead	Disables the read ahead capability.

(i) NOTE: Adaptive read ahead is no longer supported. Selecting adaptive read ahead is equivalent to selecting the read ahead option.

24 Features of PowerEdge RAID Controller 11


Virtual disk migration

The PERC 11 series supports migration of virtual disks from one controller to another without taking the target controller offline. The controller can import RAID virtual disks in optimal, degraded, or partially degraded states. You cannot import a virtual disk that is offline. When a controller detects a configured physical disk, it marks the physical disk as foreign, and generates an alert indicating that a foreign disk was detected.

Disk migration pointers:

- Supports migration of virtual disks from H740P, H745, H745P MX, and H840 to the PERC 11 series except for H345.
- Supports migration of volumes that are created within the PERC 11 series.
- Does not support migration from the PERC 11 series to PERC H345, H740P, H745, H745P MX, and H840.
- Does not support migration from PERC H330, H730, and H830 to the PERC 11 series.

(i) NOTE: The source controller must be offline before performing the disk migration.

(i) NOTE: Importing non–RAID drives and uneven span RAID 10 virtual disks from PERC 9 to PERC 11 is not supported.

(i) NOTE: Disks cannot be migrated to older generations of PERC cards.

- **NOTE:** Importing secured virtual disks is supported as long as the appropriate local key management (LKM) is supplied or configured.
- NOTE: Virtual disk migration from PERC H755 adapter, PERC H755 front SAS, PERC H755N front NVMe, PERC H750 adapter SAS, PERC H755 MX adapter to PERC H350 adapter SAS, PERC H350 Mini Monolithic SAS, PERC H355 front SAS, and PERC H355 adapter SAS is not supported.

CAUTION: Do not attempt disk migration during RLM or online capacity expansion (OCE), this causes loss of the virtual disk.

Virtual disk initialization

PERC 11 series controllers support two types of virtual disk initialization:

- Full initialization
- Fast initialization

CAUTION: Initializing virtual disks erases files and file systems while keeping the virtual disk configuration intact.

Full initialization

Performing a full initialization on a virtual disk overwrites all blocks and destroys any data that previously existed on the virtual disk. Full initialization of a virtual disk eliminates the need for the virtual disk to undergo a Background Initialization (BGI). Full initialization can be performed after the virtual disk is created.

You can start a full initialization on a virtual disk by using the Slow Initialize option in the Dell OpenManage storage management application. For more information on using the HII Configuration Utility to perform a full initialization, see Configure virtual disk parameters.

(i) NOTE: If the system reboots during a full initialization, the operation aborts and a BGI begins on the virtual disk.

Fast initialization

A fast initialization on a virtual disk overwrites the first and last 8 MB of the virtual disk, clearing any boot records or partition information. The operation takes only 2–3 seconds to complete, but it is followed by BGI, which takes a longer time to complete. To perform a fast initialization using the HII Configuration Utility, see Configure virtual disk parameters.

NOTE: During full or fast initialization, the host cannot access the virtual disk. As a result, if the host attempts to access the virtual disk while it is initializing, all I/O sent by the host will fail.



NOTE: When using iDRAC to create a virtual disk, the drive undergoes fast initialization. During this process all I/O requests to the drive will respond with a sense key of **"Not Ready"** and the I/O operation will fail. If the operating system attempts to read from the drive as soon as it discovers the drive, and while the fast initialization is still in process, then the I/O operation fails and the operating system reports an I/O error.

Reconfigure virtual disks

An online virtual disk can be reconfigured in ways that expands its capacity and changes its RAID level.

Where: Spanned virtual disks such as RAID 50 and 60 cannot be reconfigured.

(i) NOTE: Reconfiguring virtual disks typically impacts disk performance until the reconfiguration operation is complete.

Online Capacity Expansion (OCE) can be done in following ways:

 If there is a single virtual disk in a disk group and free space is available, the capacity of a virtual disk can be expanded within that free space. If multiple virtual disks exist within a common disk group, the capacities of those virtual disks cannot be expanded.

NOTE: Online capacity expansion is allowed on a disk group with a single virtual disk that begins at the start of the physical disk. It is not allowed when there is a free space at the beginning of a disk.

- 2. Add additional physical disks to a virtual disk to increase its capacity.
- **3.** After replacing all array members with larger drives than the original members, use the PERC CLI utility to expand the existing virtual disk to a larger size using the expandarray parameter. For more information, see Dell PowerEdge RAID Controller Command Line Interface Reference Guide.

RAID level migration (RLM) refers to changing a virtual disk's RAID level. Both RLM and OCE can be done simultaneously so that a virtual disk can simultaneously have its RAID level that is changed and its capacity increased. When an RLM or an OCE operation is complete, a reboot is not required.

CAUTION: Do not attempt disk migration during RLM or OCE operations. This causes loss of the virtual disk.

- **NOTE:** If an RLM or an OCE operation is in progress, then an automatic drive rebuild or copyback operation will not start until the operation is complete.
- **NOTE:** If the controller already contains the maximum number of virtual disks, you cannot perform a RAID level migration or capacity expansion on any virtual disk.
- **NOTE:** The controller changes the write cache policy of all virtual disks to write-through until the RLM or OCE operation is complete.
- (i) NOTE: You cannot initiate an OCE or an RLM on any virtual disk on the controller where a virtual disk with an ID of 0 exists.

See the following table for a list of RLM or OCE options: The source RAID level column indicates the virtual disk RAID level before the RLM or OCE operation and the target RAID level column indicates the RAID level after the RLM or OCE operation.

Table 7. RAID level migration

Source RAID Level	Target RAID Level	Number of Physical Disks (Beginning)	Number of Physical Disks (End)	Capacity Expansio n Possible	Description
RAID 0	RAID 0	1 or more	2 or more	Yes	Increases capacity by adding disks.
RAID 0	RAID 1	1	2	Yes	Converts a non-redundant virtual disk into a mirrored virtual disk by adding one disk.
RAID 0	RAID 5	1 or more	3 or more	Yes	Adds distributed parity redundancy; at least one disk must be added.



Table 7. RAID level migration (continued)

Source RAID Level	Target RAID Level	Number of Physical Disks (Beginning)	Number of Physical Disks (End)	Capacity Expansio n Possible	Description
RAID 0	RAID 6	1 or more	4 or more	Yes	Adds dual distributed parity redundancy; at least two disks must be added.
RAID 1	RAID 0	2	2 or more	Yes	Removes redundancy while increasing capacity.
RAID 1	RAID 5	2	3 or more	Yes	Maintains redundancy while adding capacity.
RAID 1	RAID 6	2	4 or more	Yes	Adds dual distributed parity redundancy and adds capacity.
RAID 5	RAID 0	3 or more	2 or more	Yes	Converts to a non-redundant virtual disk and reclaims disk space that is used for distributed parity data; one disk can be removed.
RAID 5	RAID 5	3 or more	4 or more	Yes	Increases capacity by adding disks.
RAID 5	RAID 6	3 or more	4 or more	Yes	Adds dual distributed parity redundancy; at least one disk needs to be added.
RAID 6	RAID 0	4 or more	2 or more	Yes	Converts to a non-redundant virtual disk and reclaims disk space that is used for distributed parity data; two disks can be removed.
RAID 6	RAID 5	4 or more	3 or more	Yes	Removes one set of parity data and reclaims disk space used for it; one disk can be removed.
RAID 6	RAID 6	4 or more	5 or more	Yes	Increases capacity by adding disks.
RAID 10	RAID 10	4 or more	6 or more	Yes	Increases capacity by adding disks; an even number of disks must be added.

(i) NOTE: You cannot perform a RAID level migration and expansion on RAID levels 50 and 60.

Background operations

Background initialization

Background initialization (BGI) is an automated process that writes parity or mirror data on newly created virtual disks. BGI does not run on RAID 0 virtual disks. You can control the BGI rate in the Dell OpenManage storage management application. Any change to the BGI rate does not take effect until the next BGI is performed.

() NOTE:

- You cannot disable BGI permanently. If you cancel BGI, it automatically restarts within five minutes.
- Unlike full or fast initialization of virtual disks, background initialization does not clear data from the physical disks.
- Consistency Check (CC) and BGI typically cause some loss in performance until the operation completes.



 PERC H355 adapter SAS, PERC H355 front SAS, PERC H350 adapter SAS, and PERC H350 Mini Monolithic SAS background operations will not run until the operating system boots.

Consistency check and BGI perform similar functions in that they both correct parity errors. However, CC reports data inconsistencies through an event notification, while BGI does not. You can start CC manually, but not BGI.

Consistency checks

Consistency Check (CC) is a background operation that verifies and corrects the mirror or parity data for fault tolerant virtual disks. It is recommended that you periodically run a consistency check on virtual disks.

You can manually start a CC using the HII Configuration Utility or the Dell OpenManage storage management application. You can schedule a CC to run on virtual disks using the Dell OpenManage storage management application. To start a CC using the HII Configuration Utility, see Perform consistency check.

(i) NOTE: CC or BGI typically causes some loss in performance until the operation completes.

CC and BGI both correct parity errors. However, CC reports data inconsistencies through an event notification, while BGI does not. You can start CC manually, but not BGI.

Hard drive features

This section lists the following hard drive features supported on PERC 11 cards in detail:

- Self-Encrypting Disks (SED)
- Instant Secure Erase (ISE)
- 4 KB sector disk drives

Self-Encrypting Disks

Select PERC 11 cards support self-encrypting disks (SEDs) for protection of data against loss or theft of SEDs. For information about cards supported, see Technical specifications. Protection is achieved by the use of encryption technology on the drives. There is one security key per controller. You can manage the security key using local key management (LKM) or OpenManage Secure Enterprise Key Manager also referred as Secure Enterprise Key Manager (SEKM). The controller use the security key to lock and unlock access to encryption-capable physical disks. To take advantage of this feature, you must:

- Have SEDs in your system, and
- Create a security key.

PERC cannot use SEDs that are secured by a non-PERC entity. Ensure that the SED is reprovisioned in an applicable manner by that non-PERC entity before connecting to PERC.

For more information, see the Security key and RAID management section.

(i) NOTE: You cannot enable security on non-optimal virtual disks.

NOTE: PERC 11 supports Trusted Computing Group Enterprise (TCG) Security Subsystem Classes (SSC) SAS or SATA SED drives and TCG Opal SSC NVMe drives.

Instant secure erase

Instant Secure Erase (ISE) drives use the same encryption technology as SED drives but do not allow the encryption key to be secured. The encryption technology allows the drive to be re-purposed and securely erased using the cryptographic erase function.

(i) NOTE: ISE drives do not provide protection against theft.

4 KB sector disk drives

PERC 11 controllers support 4 KB sector disk drives, which enables you to efficiently use the storage space.

28 Features of PowerEdge RAID Controller 11



Before installing Windows on 4 KB sector disk drives, see Windows operating system installation errors.

NOTE: Mixing 512–byte native and 512–byte emulated drives in a virtual disk is allowed, but mixing 512–byte and 4 KB native drives in a virtual disk is not allowed.

(i) NOTE: 4 K is only supported in UEFI mode and not legacy BIOS.

(i) NOTE: 4 K devices do not appear under the select boot device option. For more information, see Enable boot support.

Fault tolerance

The PERC 11 series supports the following:

- Self-Monitoring and Reporting Technology (SMART)
- Patrol read
- Physical disk failure detection
- Controller cache
- Battery Transparent Learn Cycle

The next sections describe some methods to achieve fault tolerance.

The SMART feature

The SMART feature monitors certain physical aspects of all motors, heads, and physical disk electronics to help detect predictable hard drive failures. Data on SMART compliant hard drives can be monitored to identify changes in values and determine whether the values are within threshold limits. Many mechanical and electrical failures display some degradation in performance before failure.

A SMART failure is also referred to as predicted failure. There are numerous factors that are predicted physical disk failures, such as a bearing failure, a broken read/write head, and changes in spin-up rate. In addition, there are factors that are related to read/write surface failure, such as seek error rate and excessive bad sectors.

NOTE: For detailed information about SCSI interface specifications, see t10.org and for detailed information about SATA interface specifications, see t13.org.

NOTE: PERC H355 adapter SAS, PERC H355 front SAS, PERC H350 adapter SAS, and PERC H350 Mini Monolithic SAS controllers do not monitor predictive failures for non-RAID disks.

Automatic Replace Member with predicted failure

A replace member operation can occur when there is a SMART predictive failure reporting on a physical disk in a virtual disk. The automatic replace member is initiated when the first SMART error occurs on a physical disk that is part of a virtual disk. The target disk needs to be a hot spare that qualifies as a rebuild disk. The physical disk with the SMART error is marked as failed only after the successful completion of the replace member. This prevents the array from reaching degraded state.

If an automatic replace member occurs using a source disk that was originally a hot spare (that was used in a rebuild), and a new disk is added and set as a target disk for the replace member operation, the hot spare drive will revert to the hot spare state after the replace member operation successfully completes.

(i) NOTE: To enable automatic replace member, use the Dell storage management application.

Patrol Read

The Patrol read feature is designed as a preventative measure to ensure physical disk health and data integrity. Patrol read scans and resolves potential problems on configured physical disks. The Dell storage management applications can be used to start patrol read and change its behavior.

The following is an overview of patrol read behavior:

- Patrol read runs on all disks on the controller that are configured as part of a virtual disk, including hot spares.
- Patrol read does not run on physical disks that are not part of a virtual disk or are in Ready state.



- The amount of controller resources dedicated to patrol read operations adjusts based on the number of outstanding disk I/O operations. For example, if the system is processing a large number of I/O operations, then patrol read uses fewer resources to allow the I/O to take a higher priority.
- Patrol read does not run on disks that are involved in any of the following operations:
 - o Rebuild
 - Replace member
 - Full or background initialization
 - CC
 - RLM or OCE

(i) NOTE: By default, patrol read automatically runs every seven days on configured SAS and SATA hard drives.

For more information about patrol read, see the Dell OpenManage documentation at OpenManage Manuals.

Physical disk failure detection

If a disk fails and it is replaced with a new disk, the controller will automatically start a rebuild on the new disk. See, Configured slot behavior. Automatic rebuilds can also occur with hot spares. If you have configured hot spares, the controller will automatically try to use them to rebuild the degraded virtual disk.

Using persistent hot spare slots

(i) NOTE: The persistent hot spare slot feature is disabled by default.

The PERC 11 series can be configured so that the system backplane or storage enclosure disk slots are dedicated as hot spare slots. This feature can be enabled using the Dell storage management application.

Once enabled, any slots with hot spares configured automatically become persistent hot spare slots. If a hot spare disk fails or is removed, a replacement disk that is inserted into the same slot automatically becomes a hot spare with the same properties as the one it is replacing. If the replacement disk does not match the disk protocol and technology, it does not become a hot spare.

For more information on persistent hot spares, see the Dell OpenManage documentation at OpenManage Manuals.

Configured slot behavior

This feature is similar to persistent hot spare slot behavior. If a redundant VD is configured to the system and if a drive is replaced, the configured slot will automatically rebuild or copyback on the inserted drive regardless of the data on the drive. This operation will overwrite the data on the drive.

Table 8. Drive state/operation

Drive state/operation	Unconfigured slot	Slot configured in VD	
Insert unconfigured drive into the system	Ready	Rebuild or copyback start	
Insert configured drive into the system	Foreign	Rebuild or copyback startOriginal drive data lost	
Insert configured locked drive into the system (unlockable)	Foreign	Cryptographic Erase (If configured VD is not secured) • Rebuild or copyback start • Original drive data lost	
Insert locked drive into the system (non-unlockable)	Foreign locked	Foreign locked	

Physical disk hot swapping

Hot swapping is the manual replacement of a disk while the PERC 11 series cards are online and performing their normal functions. The following requirements must be met before hot swapping a physical disk:

Inserido ao protocolo **22.951.206-4** por: **Pedro Henrique Golin Linhares** em: 31/03/2025 12:49. A autenticidade deste documento pode ser validada no endereço: https://www.eprotocolo.pr.gov.br/spiweb/validarDocumento com o código: 692a5ef82beb173ed33dad4ee6505546.



- The system backplane or enclosure must support hot swapping for the PERC 11 series cards.
- The replacement disk must be of the same protocol and disk technology. For example, only a SAS hard drive can replace a SAS hard drive and only a NVMe drive can replace a NVMe drive.

Using replace member and revertible hot spares

The replace member functionality allows a previously commissioned hot spare to revert to a usable hot spare. When a disk failure occurs within a virtual disk, an assigned hot spare, dedicated, or global, is commissioned and begins rebuilding until the virtual disk is optimal. After the failed disk is replaced in the same slot and the rebuild to the hot spare is complete, the controller automatically starts to copy data from the commissioned hot spare to the newly inserted disk. After the data is copied, the new disk is a part of the virtual disk and the hot spare is reverted to being a ready hot spare. This allows hot spares to remain in specific enclosure slots. While the controller is reverting the hot spare, the virtual disk remains optimal. The controller automatically reverts a hot spare only if the failed disk is replaced with a new disk in the same slot. If the new disk is not placed in the same slot, a manual replace member operation can be used to revert a previously commissioned hot spare.

NOTE: A replace member operation typically causes a temporary impact to disk performance. Once the operation completes, performance returns to normal.

Controller cache

The PERC 11 series of cards contain local DRAM on the controllers. This DRAM can cache I/O operations for Write Back, Read Ahead virtual disks to improve the performance.

NOTE: Virtual disks consisting of SSDs may not see a difference in performance using controller cache and may benefit by Fastpath.

I/O workload that is slow to HDDs, such as random 512 B and 4 kB, may take some time to flush cached data. Cache is flushed periodically but for configuration changes or system shutdown, the cache is required to be flushed before the operation can be completed. It can take several minutes to flush cache for some workloads depending on the speed of the HDDs and the amount of data in the cache.

The following operations require a complete cache flush:

- Configuration changes (add or delete VDs, VD cache setting changes, foreign configuration scan, and import)
- System reboot or shutdown
- Abrupt power loss causing cache preservation
- () NOTE: The iDRAC or OpenManage periodically scans for the foreign configurations when the foreign disks are present. This action degrades the performance. If a foreign disk is present, it is recommended that you import, clear, or remove the foreign disk to prevent an impact on the performance.

Controller cache preservation

The controller is capable of preserving its cache in the event of a system power outage or improper system shutdown. The PERC 11 series controller is attached to a battery backup unit (BBU) that provides backup power during system power loss to preserve the controller's cache data.

Cache preservation with non-volatile cache

The non-volatile cache (NVC) allows controller cache data to be stored indefinitely. If the controller has data in the cache memory during a power outage or improper system shutdown, a small amount of power from the battery is used to transfer the cache data to non-volatile flash storage where it remains until power is restored and the system is booted. If the cache preservation process is interrupted by power-on, the controller may request an extra reset during the boot to complete the process. The system displays a message during boot as Dell PERC at Bus <X> Dev <Y> has requested a system reset. System will reboot in 5 seconds.



Recovering cache data

About this task

Complete these steps if a system power loss or improper system shutdown has occurred.

Steps

- 1. Restore the system power.
- **2.** Boot the system.
- **3.** When preserved cache exists on the controller, an error message is shown. For more information about how to recover cache, see Preserved Cache State.

Battery Transparent Learn Cycle

A transparent learn cycle is a periodic operation that calculates the charge that is remaining in the battery to ensure that there is sufficient energy. The operation runs automatically, and causes no impact to the system or controller performance.

The controller automatically performs the transparent learn cycle (TLC) on the battery to calibrate and gauge its charge capacity once every 90 days. The operation can be performed manually if required.

NOTE: Virtual disks stay in write-back mode, if enabled, during transparent learn cycle. When the TLC completes, the controller sets the next TLC to +90 days.

Transparent Learn Cycle completion time

The time frame for completion of a learn cycle is a function of the battery charge capacity and the discharge and charge currents used. Typical time completion for a transparent learn cycle is between 4 to 8 hours. If the learn cycle is interrupted mid cycle, it begins at a new cycle.

Conditions for replacing the battery

The PERC battery is marked failed when the state or health of the battery is declared bad. If the battery is declared failed, then all the virtual disks in write-back mode transitions to write-through mode, and the firmware runs learn cycles in subsequent reboots until the battery is replaced. On replacing the battery, virtual disk transitions to write-back mode.

Linux operating system device enumeration

Virtual disks and non-RAID disks are presented to the operating system as SCSI devices. The operating system enumerates these devices based on the SCSI target device ID.

Enumeration order for PERC H355 adapter SAS, PERC H355 front SAS, PERC H350 adapter SAS, and PERC H350 Mini Monolithic SAS

Steps

- 1. Non-RAID disks are enumerated first.
- Virtual disks (VDs) are enumerated second, based on virtual disk target ID. Target IDs are assigned to the VDs in the ascending order when they are created. The first created VD is assigned the lowest available target ID, and the last created VD is assigned the highest available target ID. The first created VD is discovered first by the operating system.
 NOTE: The PERC H355 adapter SAS, PERC H355 front SAS, PERC H350 adapter SAS, and PERC H350 Mini

Monolithic SAS non-RAID disks may not appear in the slot order.



Enumeration order for PERC H755 front SAS, PERC H755 front SAS, PERC H755N front NVMe, PERC H750 adapter SAS, and PERC H755 MX adapter

This section describes the order of enumerating PERC H-series controlers.

Steps

- 1. Non-RAID disks are enumerated first based on slot ID.
- 2. Virtual disks (VDs) are enumerated, second based on the virtual disk target ID. Target IDs are assigned to the VDs in the descending order when they are created. The first created VD is assigned the highest available target ID, and the last created VD is assigned the lowest available target ID. Therefore, the last created VD is discovered first by the operating system.
 - () **NOTE:** Operating system enumeration may not be in this order if virtual disks or non-RAID disks are created while the operating system is running. The operating system may name devices based on the order in which they were created resulting in the operating system enumeration changing after reboot. It is recommended to reboot the system for the final device enumeration after creating any virtual disks or non-RAID disks.



Install and remove a PERC 11 card

Topics:

- Safety instructions
- Before working inside your system
- Remove the PERC H755 adapter
- Install the PERC H755 adapter
- Remove the PERC H755 front SAS card
- Install the PERC H755 front SAS card
- Remove the PERC H755N front NVMe card
- Install the PERC H755N front NVMe card
- Remove the PERC H755 MX adapter
- Install the PERC H755 MX adapter
- Remove the PERC H750 adapter SAS
- Install the PERC H750 adapter SAS
- Remove the PERC H355 adapter SAS
- Install the PERC H355 adapter SAS
- Remove the PERC H355 front SAS
- Install the PERC H355 front SAS card
- Remove the PERC H350 adapter SAS
- Install the PERC H350 adapter SAS
- Remove PERC H350 Mini Monolithic SAS
- Install PERC H350 Mini Monolithic SAS

Safety instructions

CAUTION: Ensure that two or more people lift the system horizontally from the box and place it on a flat surface, rack lift, or into the rails. WARNING: Opening or removing the PowerEdge server cover while the server is powered on may expose you to a risk of electric shock. WARNING: Do not operate the server without the cover for a duration exceeding five minutes. Operating the system without the system cover can result in component damage. NOTE: Many repairs may only be done by a certified service technician. You should only perform troubleshooting and simple repairs as authorized in your product documentation, or as directed by the online or telephone service and support team. Damage due to servicing that is not authorized by Dell is not covered by your warranty. Read and follow the safety instructions that are shipped with your product. CAUTION: To ensure proper operation and cooling, all system bays and fans must always be populated with a component or a blank. NOTE: It is recommended that you always use an antistatic mat and antistatic strap while working on components inside the server. () NOTE: To ensure proper operation and cooling, all system bays and fans must always be populated with a component or a blank. () NOTE: While replacing the hot swappable PSU, after next server boot, the new PSU automatically updates to the same firmware and configuration of the replaced one.



Before working inside your system

Prerequisites

Follows the steps listed in Safety instructions.

Steps

- 1. Power off the system and all attached peripherals.
- 2. Disconnect the system from the electrical outlet and disconnect the peripherals.
- **3.** If applicable, remove the system from the rack.

For more information, see the Rail Installation Guide relevant to your rail solutions at PowerEdge Manuals.

4. Remove the system cover.

Remove the PERC H755 adapter

Prerequisites

CAUTION: Many repairs may only be done by a certified service technician. You should only perform troubleshooting and simple repairs as authorized in your product documentation, or as directed by the online or telephone service and support team. Damage due to servicing that is not authorized by Dell is not covered by your warranty. Read and follow the safety instructions that are shipped with your product.

NOTE: It is recommended that you always use a static mat and static strap while working on components in the interior of the system.

Steps

1. Turn off the system, including any attached peripherals, and disconnect the system from the electrical outlet and peripherals.

NOTE: Perform a graceful shutdown of the system to ensure data in the cache is moved to the disk before the controller is removed.

- 2. Open the system.
- 3. Locate the PERC card in the expansion riser on the system board.

CAUTION: To prevent damage to the card, you must hold the card by its edges only.

- 4. Unfasten and lift the riser from the system board. Remove the PERC card.
- 5. Disconnect any cables connected to the card:
 - a. Press down and hold the metal tab on the cable connector.
 - **b.** Pull the cable out of the connector.
- 6. Replace the storage controller card and reconnect the data cables before placing them in the riser. For more information on installing the card, see Install PERC H755 adapter.
- 7. Reinstall the riser on the system board and fasten the riser.
- 8. Close the system.
- 9. Reconnect the system to its electrical outlet and turn the system on, including any attached peripherals.





Figure 10. Remove the PERC H755 adapter

Install the PERC H755 adapter

Prerequisites

CAUTION: Many repairs may only be done by a certified service technician. You should only perform troubleshooting and simple repairs as authorized in your product documentation, or as directed by the online or telephone service and support team. Damage due to servicing that is not authorized by Dell is not covered by your warranty. Read and follow the safety instructions that are shipped with your product.

NOTE: It is recommended that you always use a static mat and static strap while working on components in the interior of the system.

Steps

- 1. Turn off the system, including any attached peripherals, and disconnect the system from the electrical outlet.
- 2. Open the system.
- **3.** Align the card-edge connector with the connector on the system board.

\triangle CAUTION: To prevent damage to the card, you must hold the card by its edges only.

- 4. Press the card-edge down until the card is fully seated.
- 5. Connect the data cable connectors to the card.
- 6. Route the data cable through the channel on the inner side of the chassis to the backplane.
- 7. Attach the connector to the corresponding connector on the backplane as labeled on the controller.
- 8. Close the system.
- 9. Reconnect the system to its electrical outlet and turn the system on, including any attached peripherals.





Figure 11. Install the PERC H755 adapter

Remove the PERC H755 front SAS card

Prerequisites

CAUTION: Many repairs may only be done by a certified service technician. You should only perform troubleshooting and simple repairs as authorized in your product documentation, or as directed by the online or telephone service and support team. Damage due to servicing that is not authorized by Dell is not covered by your warranty. Read and follow the safety instructions that are shipped with your product.

NOTE: It is recommended that you always use a static mat and static strap while working on components in the interior of the system.

Steps

1. Turn off the system, including any attached peripherals, and disconnect the system from the electrical outlet and peripherals.

NOTE: Perform a graceful shutdown of the system to ensure data in the cache is moved to the disk before the controller is removed.

- 2. Open the system.
- 3. Locate the PERC card in the controller carrier at the front of the system.

CAUTION: To prevent damage to the card, you must hold the card by its edges only.

4. Unscrew the fasteners on the controller carrier and slide the carrier away from the backplane, disconnecting the controller from the backplane.

If you are removing a PERC H755 front SAS controller in the upside down orientation, you must remove both the backplane and the controller at the same time because of the limited clearance available:

- a. Uninstall all drives from the backplane.
- b. Disconnect all cables between the PERC and the backplane.
- c. Lift the backplane and PERC from the system.
- 5. Disconnect any cables connected to the card:
 - a. Press down and hold the metal tab on the cable connector.
 - ${\bf b.}~$ Pull the cables out of the connector.
- 6. Remove the PERC controller from the controller carrier.

Install and remove a PERC 11 card 37



- 7. Insert the replacement controller into the carrier and secure it with the appropriate screws.
- 8. Take the replacement storage controller and reconnect the cables before reconnecting it to the backplane. If you are removing a PERC H755 front SAS controller in the upside down orientation, reattach the PERC controller to the backplane first before reinstalling the backplane into the system. For more information on installing the card, see Install PERC H755 front SAS card.
- 9. Close the system.
- 10. Reconnect the system to its electrical outlet and turn the system on, including any attached peripherals.



Figure 12. Remove the PERC H755 front SAS card

Install the PERC H755 front SAS card

Prerequisites

- CAUTION: Many repairs may only be done by a certified service technician. You should only perform troubleshooting and simple repairs as authorized in your product documentation, or as directed by the online or telephone service and support team. Damage due to servicing that is not authorized by Dell is not covered by your warranty. Read and follow the safety instructions that are shipped with your product.
- **NOTE:** It is recommended that you always use a static mat and static strap while working on components in the interior of the system.

Steps

1. Turn off the system, including any attached peripherals, and disconnect the system from the electrical outlet.



NOTE: Perform a graceful shutdown of the sled to ensure that data in the cache is moved to the disk before the controller is removed.

- 2. Open the system.
- 3. Connect the PERC card to the carrier and ensure that the screws are properly fastened in place.

CAUTION: To prevent damage to the card, hold the card by its edges only.

- **4.** Align the carrier with the guide pins until the controller is securely seated.
- 5. Slide the card into the connector until it is fully seated in the connector. Tighten the screws on the carrier that connect to the chassis to secure the carrier.
- 6. Connect the cable connectors to the card.

NOTE: Ensure that you connect the cable according to the connector labels on the cable. The cable does not function properly if reversed.

- 7. Close the system.
- 8. Reconnect the system to its electrical outlet and turn on the system and any attached peripherals.



Figure 13. Install the PERC H755 front SAS card

Remove the PERC H755N front NVMe card

Prerequisites

CAUTION: Many repairs may only be done by a certified service technician. You should only perform troubleshooting and simple repairs as authorized in your product documentation, or as directed by the online or

Install and remove a PERC 11 card 39



telephone service and support team. Damage due to servicing that is not authorized by Dell is not covered by your warranty. Read and follow the safety instructions that are shipped with your product.

NOTE: It is recommended that you always use a static mat and static strap while working on components in the interior of the system.

Steps

- 1. Turn off the system, including any attached peripherals, and disconnect the system from the electrical outlet and peripherals.
 - **NOTE:** Perform a graceful shutdown of the system to ensure that data in the cache is moved to the disk before the controller is removed.
- 2. Open the system.
- 3. Locate the PERC card in the controller carrier at the front of the system.

CAUTION: To prevent damage to the card, you must hold the card by its edges only.

4. Unscrew the fasteners on the controller carrier, and slide the carrier away from the backplane to disconnect the controller from the backplane.

If you are removing a PERC H755N front NVMe controller in the upside down orientation, you must remove both the backplane and the controller at the same time because of the limited clearance available:

- **a.** Uninstall all drives from the backplane.
- **b.** Disconnect all cables between the PERC and the backplane.
- c. Lift the backplane and PERC from the system.
- 5. Disconnect any cables connected to the card:
 - a. Press down and hold the metal tab on the cable connector.
 - **b.** Pull the cable out of the connector.
- 6. Remove the PERC controller from the controller carrier.
- 7. Insert the replacement controller into the carrier and secure it with the appropriate screws.
- 8. Take the replacement storage controller and reconnect the cable before reconnecting it to the backplane.

If you are removing a PERC H755 front NVMe controller in the upside down orientation, reattach the PERC controller to the backplane first before reinstalling the backplane into the system.

- 9. Close the system.
- 10. Reconnect the system to its electrical outlet and turn the system on, including any attached peripherals.





Figure 14. Remove the PERC H755N front NVMe card

Install the PERC H755N front NVMe card

Prerequisites

- CAUTION: Many repairs may only be done by a certified service technician. You should only perform troubleshooting and simple repairs as authorized in your product documentation, or as directed by the online or telephone service and support team. Damage due to servicing that is not authorized by Dell is not covered by your warranty. Read and follow the safety instructions that are shipped with your product.
- **NOTE:** It is recommended that you always use a static mat and static strap while working on components in the interior of the system.

Steps

- 1. Turn off the system, including any attached peripherals, and disconnect the system from the electrical outlet.
 - () NOTE: Perform a graceful shutdown of the sled to ensure that data in the cache is moved to the disk before the controller is removed.
- 2. Open the system.
- 3. Connect the PERC card to the carrier and ensure the screws are properly fastened in place.

CAUTION: To prevent damage to the card, hold the card by its edges only.

4. Align the carrier with the guide pins until the controller is securely seated.



- 5. Slide the card until it is fully seated in the connector. Tighten the screws on the carrier that connect to the chassis to secure the carrier.
- ${\bf 6.}\$ Connect the cable connectors to the card.

() NOTE: Ensure that you connect the cable according to the connector labels on the cable. The cable does not function properly if reversed.

- 7. Close the system.
- 8. Reconnect the system to its electrical outlet and turn on the system and any attached peripherals.



Figure 15. Install the PERC H755N front NVMe card

Remove the PERC H755 MX adapter

Prerequisites

CAUTION: Many repairs may only be done by a certified service technician. You should only perform troubleshooting and simple repairs as authorized in your product documentation, or as directed by the online or telephone service and support team. Damage due to servicing that is not authorized by Dell is not covered by your warranty. Read and follow the safety instructions that are shipped with your product.

CAUTION: To prevent damage to the card, hold the card by its edges only.

() NOTE: It is recommended that you always use a static mat and static strap while working on components in the interior of the system.



Steps

- 1. Turn off the sled, including any attached peripherals, and remove the sled from the MX chassis.
 - **NOTE:** Perform a graceful shutdown of the system to ensure that data in the cache is moved to the disk before the controller is removed.
- 2. Open the sled.
- **3.** Locate the PERC card on the system board.

CAUTION: To prevent damage to the card, hold the card by its edges only.

- 4. Using the blue tab, rotate the lever of the controller.
- 5. Pull the release lever upward to disengage the controller from the connector.
- 6. Disconnect the cable from the card. To disconnect the cable:
 - **a.** Press and hold the metal tab on the cable connector.
 - **b.** Pull the cable out of the connector.
- 7. Lift the card from the system board.
- 8. Replace the storage controller card and connect the cable. For information on installing the card, see Install the PERC H755 MX adapter.
- 9. Close the sled.
- 10. Insert the sled into the MX chassis and turn on the system and any attached MX chassis peripherals.



Figure 16. Remove the PERC H755 MX adapter

Install the PERC H755 MX adapter

Prerequisites

CAUTION: Many repairs may only be done by a certified service technician. You should only perform troubleshooting and simple repairs as authorized in your product documentation, or as directed by the online or



telephone service and support team. Damage due to servicing that is not authorized by Dell is not covered by your warranty. Read and follow the safety instructions that are shipped with your product.

NOTE: It is recommended that you always use a static mat and static strap while working on components in the interior of the system.

Steps

- 1. Turn off the sled and any attached peripherals, and remove the sled from the MX chassis.
- 2. Open the sled.
- **3.** Connect the backplane data cable connector to the card.
 - **NOTE:** Ensure that you connect the cable according to the connector labels on the cable. The cable does not function properly if reversed.
- **4.** Align the bracket notches with the tabs on the sides of the sled chassis and align the PERC card connector with the connector on the system board.

CAUTION: To prevent damage to the card, hold the card by its edges only.

- 5. Press the PERC card into the connector until it is firmly seated.
- 6. Press the release lever to secure the card to the sled.

(i) NOTE: The pin on the release lever secures the card to the chassis of the sled.

- 7. Route the data cable through the clip on the card and through the channel on the inner side of the chassis.
- 8. Attach the connector to the corresponding connector on the backplane as labeled in the controller.
- 9. Close the sled.
- 10. Insert the sled into the MX chassis and turn on the system and any attached MX chassis peripherals.



Figure 17. Install the PERC H755 MX adapter



Remove the PERC H750 adapter SAS

Prerequisites

CAUTION: Many repairs may only be done by a certified service technician. You should only perform troubleshooting and simple repairs as authorized in your product documentation, or as directed by the online or telephone service and support team. Damage due to servicing that is not authorized by Dell is not covered by your warranty. Read and follow the safety instructions that are shipped with your product.

() NOTE: It is recommended that you always use a static mat and static strap while working on components in the interior of the system.

Steps

- 1. Turn off the system, including any attached peripherals, and disconnect the system from the electrical outlet and peripherals.
- 2. Open the system.
- 3. Locate the PERC card on the system board.

CAUTION: To prevent damage to the card, you must hold the card by its edges only.

- 4. Lift the card to remove it from the connector on the system board.
- 5. Disconnect the SAS cables connected to the card:
 - a. Press down and hold the metal tab on the SAS cable connector.
 - **b.** Pull the SAS cable out of the connector.
- 6. Replace the storage controller card and connect the cable. For more information on installing the card, see Install the H750 adapter SAS.
- 7. Close the system.
- 8. Reconnect the system to its electrical outlet and turn the system on, including any attached peripherals.



Figure 18. Remove PERC H750 adapter SAS

Install the PERC H750 adapter SAS

Prerequisites

CAUTION: Many repairs may only be done by a certified service technician. You should only perform troubleshooting and simple repairs as authorized in your product documentation, or as directed by the online or telephone service and support team. Damage due to servicing that is not authorized by Dell is not covered by your warranty. Read and follow the safety instructions that are shipped with your product.

Install and remove a PERC 11 card 45



NOTE: It is recommended that you always use a static mat and static strap while working on components in the interior of the system.

Steps

- 1. Turn off the system, including any attached peripherals, and disconnect the system from the electrical outlet.
- 2. Open the system.
- 3. Align the card-edge connector with the connector on the system board.

\triangle CAUTION: To prevent damage to the card, you must hold the card by its edges only.

- 4. Press the card-edge down until the card is fully seated.
- 5. Connect the SAS data cable connector to the card.

NOTE: Ensure that you connect the cable according to the connector labels on the cable. The cable does not function properly if reversed.

- 6. Route the SAS data cable through the channel on the inner side of the chassis to the backplane.
- 7. Attach the connector labeled SAS A to connector SAS A on the backplane.
- 8. Close the system.
- 9. Reconnect the system to its electrical outlet and turn the system on, including any attached peripherals.



Figure 19. Install PERC H750 adapter SAS

Remove the PERC H355 adapter SAS

Describes the tasks to remove a PERC H355 adapter SAS controller from a server.

Prerequisites

CAUTION: Many repairs may only be done by a certified service technician. You should only perform troubleshooting and simple repairs as authorized in your product documentation, or as directed by the online or telephone service and support team. Damage due to servicing that is not authorized by Dell is not covered by your warranty. Read and follow the safety instructions that are shipped with your product.

NOTE: It is recommended that you always use a static mat and static strap while working on components in the interior of the system.

Steps

1. Turn off the system, including any attached peripherals, and disconnect the system from the electrical outlet and peripherals.



- 2. Open the system.
- 3. Locate the PERC card in the expansion riser on the system board.

CAUTION: To prevent damage to the card, you must hold the card by its edges only.

- 4. Unfasten and lift the riser from the system board. Remove the PERC card.
- **5.** Disconnect any SAS cables connected to the card:
 - a. Press down and hold the metal tab on the SAS cable connector.
 - **b.** Pull the SAS cable out of the connector.
- 6. Replace the storage controller and reconnect the SAS cable before placing them in the riser. For more information on installing the card, see Install the PERC H355 adapter.
- 7. Reinstall the riser on the system board and fasten the riser.
- 8. Close the system.
- 9. Reconnect the system to its electrical outlet and turn the system on, including any attached peripherals.



Figure 20. Remove the PERC H355 adapter SAS

Install the PERC H355 adapter SAS

Describes the tasks to install a PERC H355 adapter SAS controller in a server.

Prerequisites

- CAUTION: Many repairs may only be done by a certified service technician. You should only perform troubleshooting and simple repairs as authorized in your product documentation, or as directed by the online or telephone service and support team. Damage due to servicing that is not authorized by Dell is not covered by your warranty. Read and follow the safety instructions that are shipped with your product.
- () NOTE: It is recommended that you always use a static mat and static strap while working on components in the interior of the system.

Steps

- 1. Turn off the system, including any attached peripherals, and disconnect the system from the electrical outlet.
- 2. Open the system.

Install and remove a PERC 11 card 47



3. Align the card-edge connector with the connector on the system board.

CAUTION: To prevent damage to the card, you must hold the card by its edges only.

- 4. Press the card-edge down until the card is fully seated.
- 5. Connect the SAS data cable connectors to the card.
 - **NOTE:** Ensure that you connect the cable according to the connector labels on the cable. The cable does not function properly if reversed.
- 6. Route the SAS data cable through the channel on the inner side of the chassis to the backplane.
- 7. Attach the connector labeled SAS A to connector SAS A on the backplane, and attach the connector labeled SAS B to connector SAS B on the backplane.
- 8. Close the system.
- 9. Reconnect the system to its electrical outlet and turn the system on, including any attached peripherals.



Figure 21. Install the PERC H355 adapter SAS

Remove the PERC H355 front SAS

Prerequisites

- CAUTION: Many repairs may only be done by a certified service technician. You should only perform troubleshooting and simple repairs as authorized in your product documentation, or as directed by the online or telephone service and support team. Damage due to servicing that is not authorized by Dell is not covered by your warranty. Read and follow the safety instructions that are shipped with your product.
- **NOTE:** It is recommended that you always use a static mat and static strap while working on components in the interior of the system.

Steps

1. Turn off the system, including any attached peripherals, and disconnect the system from the electrical outlet and peripherals.

NOTE: Perform a graceful shutdown of the system to ensure data in the cache is moved to the disk before the controller is removed.



- 2. Open the system.
- 3. Locate the PERC card in the controller carrier at the front of the system.

CAUTION: To prevent damage to the card, you must hold the card by its edges only.

4. Unscrew the fasteners on the controller carrier and slide the carrier away from the backplane, disconnecting the controller from the backplane.

If you are removing a PERC H355 front SAS controller in the upside down orientation, you must remove both the backplane and the controller at the same time because of the limited clearance available:

- **a.** Uninstall all drives from the backplane.
- b. Disconnect all cables between the PERC and the backplane.
- c. Lift the backplane and PERC from the system.
- **5.** Disconnect any cables connected to the card:
 - a. Press down and hold the metal tab on the cable connector.
 - b. Pull the cables out of the connector.
- 6. Remove the PERC controller from the controller carrier.
- 7. Insert the replacement controller into the carrier and secure it with the appropriate screws.
- 8. Take the replacement storage controller and reconnect the cables before reconnecting it to the backplane.
- If you are removing a PERC H355 front SAS controller in the upside down orientation, reattach the PERC controller to the backplane first before reinstalling the backplane into the system. For more information on installing the card, see Install the PERC H355 front.
- 9. Close the system.
- **10.** Reconnect the system to its electrical outlet and turn the system on, including any attached peripherals.



Figure 22. Remove the PERC H355 front SAS



Install the PERC H355 front SAS card

Prerequisites

CAUTION: Many repairs may only be done by a certified service technician. You should only perform troubleshooting and simple repairs as authorized in your product documentation, or as directed by the online or telephone service and support team. Damage due to servicing that is not authorized by Dell is not covered by your warranty. Read and follow the safety instructions that are shipped with your product.

NOTE: It is recommended that you always use a static mat and static strap while working on components in the interior of the system.

Steps

- 1. Turn off the system, including any attached peripherals, and disconnect the system from the electrical outlet.
 - (i) **NOTE:** Perform a graceful shutdown of the sled to ensure that data in the cache is moved to the disk before the controller is removed.
- 2. Open the system.
- 3. Connect the PERC card to the carrier and ensure that the screws are properly fastened in place.

CAUTION: To prevent damage to the card, hold the card by its edges only.

- 4. Align the carrier with the guide pins until the controller is securely seated.
- 5. Slide the card into the connector until it is fully seated in the connector. Tighten the screws on the carrier that connect to the chassis to secure the carrier.
- 6. Connect the cable connectors to the card.

NOTE: Ensure that you connect the cable according to the connector labels on the cable. The cable does not function properly if reversed.

- 7. Close the system.
- 8. Reconnect the system to its electrical outlet and turn on the system and any attached peripherals.





Figure 23. Install the PERC H755 front SAS card

Remove the PERC H350 adapter SAS

Prerequisites

CAUTION: Many repairs may only be done by a certified service technician. You should only perform troubleshooting and simple repairs as authorized in your product documentation, or as directed by the online or telephone service and support team. Damage due to servicing that is not authorized by Dell is not covered by your warranty. Read and follow the safety instructions that are shipped with your product.

NOTE: It is recommended that you always use a static mat and static strap while working on components in the interior of the system.

Steps

- 1. Turn off the system, including any attached peripherals, and disconnect the system from the electrical outlet and peripherals.
- **2.** Open the system.
- **3.** Locate the PERC card on the system board.

CAUTION: To prevent damage to the card, you must hold the card by its edges only.

- 4. Lift the card to remove it from the connector on the system board.
- **5.** Disconnect the SAS cables connected to the card:
 - **a.** Press down and hold the metal tab on the SAS cable connector.
 - **b.** Pull the SAS cable out of the connector.

Install and remove a PERC 11 card 51



- 6. Replace the storage controller card and connect the cable. For more information on installing the card, see Install the PERC H350 adapter.
- 7. Close the system.
- 8. Reconnect the system to its electrical outlet and turn the system on, including any attached peripherals.



Figure 24. Remove the PERC H350 adapter SAS

Install the PERC H350 adapter SAS

Prerequisites

- CAUTION: Many repairs may only be done by a certified service technician. You should only perform troubleshooting and simple repairs as authorized in your product documentation, or as directed by the online or telephone service and support team. Damage due to servicing that is not authorized by Dell is not covered by your warranty. Read and follow the safety instructions that are shipped with your product.
- **NOTE:** It is recommended that you always use a static mat and static strap while working on components in the interior of the system.

Steps

- 1. Turn off the system, including any attached peripherals, and disconnect the system from the electrical outlet.
- 2. Open the system.
- **3.** Align the card-edge connector with the connector on the system board.

\triangle CAUTION: To prevent damage to the card, you must hold the card by its edges only.

- 4. Press the card-edge down until the card is fully seated.
- 5. Connect the SAS data cable connector to the card.

NOTE: Ensure that you connect the cable according to the connector labels on the cable. The cable does not function properly if reversed.

- 6. Route the SAS data cable through the channel on the inner side of the chassis to the backplane.
- 7. Attach the connector labeled SAS A to connector SAS A on the backplane.
- 8. Close the system.
- 9. Reconnect the system to its electrical outlet and turn the system on, including any attached peripherals.





Figure 25. Install the PERC H350 adapter SAS

Remove PERC H350 Mini Monolithic SAS

Prerequisites

CAUTION: Many repairs may only be done by a certified service technician. You should only perform troubleshooting and simple repairs as authorized in your product documentation, or as directed by the online or telephone service and support team. Damage due to servicing that is not authorized by Dell is not covered by your warranty. Read and follow the safety instructions that are shipped with your product.

Steps

- 1. Using Phillips #2 screwdriver, loosen the screws that secure the storage controller cable to the connector on the system board.
- 2. Lift the storage controller cable to disconnect it from the connector on the system board.







- 3. Lift one end of the card and angle it to disengage the card from the card holder on the system board.
- **4.** Lift the card out of the system.



Figure 27. Remove the PERC H350 Mini Monolithic SAS



Install PERC H350 Mini Monolithic SAS

Prerequisites

CAUTION: Many repairs may only be done by a certified service technician. You should only perform troubleshooting and simple repairs as authorized in your product documentation, or as directed by the online or telephone service and support team. Damage due to servicing that is not authorized by Dell is not covered by your warranty. Read and follow the safety instructions that are shipped with your product.

Steps

- 1. Angle the integrated storage controller card and align the end of the card with the storage controller card connector on the system board.
- 2. Lower the connector side of the storage controller card into the storage controller card connector on the system board.

(i) NOTE: Ensure that the slots on the system board align with the screw holes on the storage controller card connector.



Figure 28. Install PERC H350 Mini Monolithic SAS

- 3. Route the storage controller card cable along with the wall of the system.
- 4. Align the screws on the integrated storage controller card cable with the screw holes on the connector.
- **5.** Using Phillips #2 screwdriver, tighten the screws to secure the integrated storage controller card cable to the card connector on the system board.





Figure 29. Install the cable

56 Install and remove a PERC 11 card

Inserido ao protocolo **22.951.206-4** por: **Pedro Henrique Golin Linhares** em: 31/03/2025 12:49. A autenticidade deste documento pode ser validada no endereço: https://www.eprotocolo.pr.gov.br/spiweb/validarDocumento com o código: **692a5ef82beb173ed33dad4ee6505546**.



Driver support for PERC 11

The PERC 11 cards require software drivers to operate with the supported operating systems.

This chapter contains the procedures for installing the drivers for the PERC 11 cards.

NOTE: The driver for PERC 11 for VMware ESXi is packaged within the VMware ESXi ISO image that is downloaded from Dell. For more information, see the VMware documentation at Virtualization Solutions Documentation. It is not recommended to have drivers from controllers prior to PERC 11 on the same system.

The two methods for installing a driver that is discussed in this chapter are:

- Installing a driver during operating system installation: Use this method if you are performing a new installation of the operating system and want to include the drivers.
- **Updating existing drivers:** Use this method if the operating system and the HBA controllers are already installed and you want to update to the latest drivers.

Topics:

- Creating the device driver media
- Windows driver installation
- Linux driver installation
- Load the driver while installing an operating system

Creating the device driver media

Use one of the following two methods to create the device driver media:

- Downloading Drivers From The Dell Support Website
- Downloading Drivers From The Dell Systems Service And Diagnostic Tools Media

Download and save PERC 11 drivers from the support site

About this task

To download drivers from the Dell Support website:

Steps

- 1. Go to the Support Site.
- 2. Enter the Service Tag of your system in the Choose by Service Tag to get started field or select Choose from a list of all Dell products.
- **3.** Select the **System Type**, **Operating System**, and **Category** from the drop-down list. The drivers that are applicable to your selection are displayed.
- 4. Download the drivers that you require to a USB drive, CD, or DVD.
- **5.** During the operating system installation, use the media that you created to load the driver. For more information about reinstalling the operating system, see the relevant section for your operating system later in this guide.

Download and save PERC 11 drivers from the Dell Systems Service and Diagnostic Tools

About this task

To download drivers from the **Dell Systems Service and Diagnostic Tools** media:



Steps

- Insert the Dell Systems Service and Diagnostics Tools media in your system. The Welcome to Dell Service and Diagnostic Utilities screen is displayed.
- 2. Select your system model and operating system.
- 3. Click Continue.
- 4. From the list of drivers displayed, select the driver you require.
- 5. Select the self-extracting ZIP file and click $\ensuremath{\textbf{Run}}$.
- 6. Copy the driver to a CD, DVD, or USB drive.
- 7. Repeat steps 1 to 6 for all the drivers you require.

Windows driver installation

Before you install the Windows driver for PERC 11, you must first create a device driver media.

- Read the Microsoft **Getting Started** document that shipped with your operating system.
- Ensure that your system has the latest BIOS, firmware, and driver updates. If required, download the latest BIOS, firmware, and driver updates from Support Site.
- Create a device driver media using one of the methods listed below:
 - USB drive
 - CD
 - o DVD

Install PERC 11 driver while newly installing the Windows Server 2016 and later

About this task

To install the driver:

Steps

- 1. Boot the system using the Windows Server 2016, or newer media.
- 2. Follow the on-screen instructions until you reach Where do you want to install Windows Server 2016 or later window and then select Load driver.
- 3. As prompted, insert the installation media and browse to the appropriate location.
- 4. Select a PERC 11 series card from the list.
- 5. Click **Next** and continue installation.

Install PERC 11 driver on which the Windows Server 2016 is already installed and later

About this task

Perform the following steps to configure the driver for the RAID controller on which the Windows Server 2016 is already installed:

Steps

- 1. Turn off the system.
- 2. Install the new RAID controller in the system.
- For detailed instructions on installing the RAID controller in the system, see Install and remove a PERC 11 card.
- **3.** Turn on the system.
 - The Found New Hardware Wizard screen displays the detected hardware device.
- 4. Click Next



- 5. On the Locate device driver screen, select Search for a suitable driver for my device and click Next.
- 6. Browse and select the drivers from the Locate Driver Files screen.
- 7. Click Next.
 - The wizard detects and installs the appropriate device drivers for the new RAID controller.
- 8. Click Finish to complete the installation.
- 9. Reboot the system when prompted.

Update PERC 11 driver that runs on Windows Server 2016 or later

Prerequisites

(i) NOTE: Close all applications on your system before you update the driver.

Steps

- 1. Insert the media containing the driver.
- 2. Select Start > Settings > Control Panel > System. The System Properties screen is displayed.

(i) NOTE: The path to System might vary depending on the operating system family.

- **3.** Click the **Hardware** tab.
- 4. Click Device Manager.

The **Device Manager** screen is displayed.

(i) NOTE: The path to Device Manager might vary depending on the operating system family.

- 5. Expand Storage Controllers by double-clicking the entry or by clicking on the plus (+) symbol next to Storage Controllers.
- 6. Double-click the controller for which you want to update the driver.
- 7. Click the Driver tab and click Update Driver.
- The screen to update the device driver wizard is displayed.
- 8. Select Install from a list or specific location.
- 9. Click Next.
- 10. Follow the steps in the wizard and browse to the location of the driver files.
- 11. Select the INF file from the drive media.
- 12. Click Next and continue the installation steps in the wizard.
- 13. Click Finish to exit the wizard and reboot the system for the changes to take place.
 - (i) NOTE: Dell provides the Dell Update Package (DUP) to update drivers on systems running Windows Server 2016 and newer operating system. DUP is an executable application that updates drivers for specific devices. DUP supports command line interface and silent execution. For more information, see Support Site.

Linux driver installation

The driver update disk (DUD) images are created only for those operating system releases in which the native (in-box) driver is insufficient for installation. In the event that an operating system is being installed with a corresponding DUD image, see, Installing or updating the RPM driver package with KMOD support. If not, proceed with using the native device driver and then skip to the topic Installing or Updating the RPM Driver Package With KMP Support.

(i) NOTE: The driver update disk (DUD) images are created only for those operating system releases in which the native (in-box) driver is insufficient for installation. In the event that an operating system is being installed with a corresponding DUD image, follow the instructions below.

(i) NOTE: To view the complete list of boot loader options, see the installation guide of your operating system.

NOTE: If using out-of-box drivers with RHEL 7 and higher, a tainted kernel message will be displayed in the log. RedHat does not provide a mechanism to sign external drivers for RHEL.



Install or update a RPM driver package using the KMOD support

Prerequisites

(i) NOTE: This procedure is applicable for Red Hat Enterprise Linux 7.x and higher.

About this task

Perform the following steps to install the RPM package with KMOD support:

Steps

- 1. Uncompress the gzipped tarball driver release package.
- 2. Install the driver package using the command: rpm -ihv kmodmegaraid_ sas-<version>.rpm.

(i) NOTE: Use rpm -Uvh <package name> when upgrading an existing package.

- 3. If the previous device driver is in use, you must reboot the system for the updated driver to take effect.
- 4. Verify the loaded driver version by running the following command: modinfo megaraid sas.

Install or update a RPM driver package using the KMP support

Prerequisites

(i) NOTE: This procedure is applicable for SUSE Enterprise Linux 15.x.

About this task

Perform the following steps to install the RPM package with KMP support:

Steps

- 1. Uncompress the gzipped tarball driver release package.
- 2. Install the driver package using the command: rpm -ihv kmpmegaraid_ sas- <version>.rpm.

(i) NOTE: Use rpm -Uvh <package name> when updating an existing package.

- 3. If the previous device driver is in use, you must reboot the system for the updated driver to take effect.
- 4. Verify the loaded driver version by running the following command: modinfo megaraid sas.

Upgrading the Kernel

About this task

When upgrading to a new kernel, you must reinstall the DKMS-enabled driver packages. Perform the following steps to update or install the driver for a new kernel:

Steps

- 1. At a terminal window, type the following: dkms build -m <module_name> v <module version> k <kernel version> dkms install -m <module name> - v <module version> - k <kernel version>.
- 2. To check if the driver is successfully installed in the new kernel, type: dkms status. A message similar to the following is displayed: <driver name>, <driver version>, <new kernel version>: installed.
- 3. If the previous device driver is in use, you must restart the system for the updated driver to take effect.


Load the driver while installing an operating system

Steps

- 1. Perform the following operations to install driver media:
 - PERC Linux driver ISO:
 - a. Download the PERC Linux driver package from the Dell Support site.
 - b. Extract two base directories from the tar.gz package (tar.gz > tar > base directories).
 - c. Extract the ISO file that is available in the zipped disks-x directory. For example, RHEL79/disks-1/ megaraid_sas-07.719.03.00_el7.9-1.x86_64.iso.gz > megaraid_sas-07.719.03.00_el7.9-1.x86_64.iso
 - **d.** Mount the ISO to the Server, burn the ISO to a CD or DVD or copy the ISO file to a USB. The USB has to match with the ISO.
 - LC driver pack:
 - a. Install the LC driver pack.
 - **b.** Boot the life-cycle controller and go through the operating system deployment wizard.
- 2. Boot to the installer.
- 3. In the Installation screen, press E.
- **4.** Perform the following operation:
 - If the operating system is Red Hat Enterprise Linux 7 or RHEL 8, the CLI displays the syntax vmlinuz. Enter inst.dd.

For example, when you are prompted with the command vmlinuz intrd=initrd.img inst.stage2=hd:LABEL=RHEL-7.0\x20x86_64 quiet inst.dd.

• If the operating system is SLES 15, the CLI displays the syntax linuxefi.. Enter **dud=1**.

For example, when you are prompted with the command linuxefi/boot/x86_64/loader/linux splash=silent dud=1.

NOTE: Boot parameters may vary based on the operating system version. See operating system installation manuals for exact boot parameter syntax.

- 5. Attach the driver media (ISO, USB).
- **6.** Press F10 to boot to the operating system. A screen is displayed prompting you to select the driver media (USB, CD, ISO, and so on).
- 7. When prompted, select the driver media.

If applicable select the PERC driver ...megaraid_sas...

(i) NOTE: Ensure that the driver is selected with an X symbol.

- 8. The driver should be extracted or loaded.
- 9. Before proceeding or exiting the driver select menu, disconnect the driver media.

NOTE: Ensure that you disconnect the driver media so that the drivers are loaded successfully. If the installation media is deleted, reattach it.

10. Press C or exit to go to the installation.





This section provides information about downloading and installing the firmware using Dell Update Package (DUP).

Topics:

• Upgrade firmware controller using Dell Update Package (DUP)

Upgrade firmware controller using Dell Update Package (DUP)

About this task

(i) NOTE: If the Online Capacity Expansion operation is in progress then you cannot update the firmware version.

Steps

- 1. Go to the Drivers and Downloads page on the support site.
- 2. Locate your controller.
- 3. Download the DUP file.
 - a. To upgrade by using Windows or iDRAC, download the Windows executable file.
 - **b.** To upgrade using Linux, download the **.bin** file.

(i) NOTE: For VMware, firmware must be upgraded by using iDRAC or the PERC CLI.

- 4. Install the DUP by doing one of the following:
 - a. For Windows, run the executable file in the Windows environment.
 - **b.** For Linux, run the **.bin** file in the Linux environment.
 - c. For iDRAC, click System iDRAC > Maintenance > System Update, upload Windows executable, and then install.



Manage PERC 11 controllers using HII configuration utility

The Human Interface Infrastructure (HII) configuration utility is a storage management application integrated into the System BIOS < F2 >. It is used to configure and manage the controller(s), virtual disks, and physical disks. This utility is independent of the operating system.

Topics:

- Enter the PERC 11 HII configuration utility
- Exit the PERC 11 HII configuration utility
- Navigate to Dell PERC 11 configuration utility
- View the HII Configuration utility dashboard
- Configuration management
- Controller management
- Virtual disk management
- Physical disk management
- Hardware components
- Security key management in HII configuration utility

Enter the PERC 11 HII configuration utility

About this task

Perform the following steps to boot to the HII configuration utility:

Steps

- 1. Turn on the system.
- 2. While the system startup, press <F2> to enter System Setup.

3. Click Device Settings.

Device Settings screen lists all the RAID controllers in the system.

To access the management menu for the controller, use the arrow keys or the mouse.

- () NOTE: For more information in all the options, click Help that is available on the top right-hand corner of the browser screen. Help information for individual option menus can also be viewed by scrolling down on each option.
- **NOTE:** Some of the options within the HII configuration utility are not present if the controller does not support the corresponding feature. Options may also be grayed out if the feature is not applicable to the current configuration.

Exit the PERC 11 HII configuration utility

About this task

To exit the HII configuration utility, perform the following steps:

Steps

- 1. Click **Finish** at the bottom-right corner on the **System Setup Main Menu** screen. Displays a warning message to confirm your choice.
- 2. Click Yes to exit the HII configuration utility.



Navigate to Dell PERC 11 configuration utility

Steps

- 1. Enter the UEFI configuration Utility. See Enter the PERC 11 HII configuration utility. The **Device Settings** screen displays a list of NIC ports and the RAID controllers.
- To enter PERC 11 configuration utility, click the appropriate PERC controllers. The Dashboard view screen is displayed.

View the HII Configuration utility dashboard

The first screen that is displayed when you access the HII Configuration Utility is the **Dashboard View** screen. The following table provides detailed information about the options available on the **Dashboard View** screen.

Table 9. Dashboard view screen

Dashboard view options	Description
Main menu	Displays the following configuration options: • Configuration Management • Controller Management • Virtual Disk Management • Physical Disk Management • Hardware Components
Help	Provides context sensitive help message.
Properties	 Displays the following information about the controller: Status — displays the status of the controller. Backplane — displays information about the number of backplanes connected to the controller. BBU — displays information about the availability of Battery Backup Unit (BBU). Enclosure — displays information about the number of enclosures connected to the controller. Physical Disks — displays information about the number of physical disks connected to the controller. Disk Groups — displays information about the number of disk groups connected to the controller. Virtual Disks — displays information about the number of virtual disks connected to the controller.
View server profile	 Displays HII Spec version supported on the system and also displays the following menu options for controller components: Controller Management Hardware Components Physical Disk Management Virtual Disk Management
Actions	 Displays the following options: Configure — displays configuration options that are supported by the controller. Set Factory Defaults — restore factory default values for all controller properties.
Background operations	Displays if virtual disk or physical disk operations are in progress.



Configuration management

Auto Configure RAID 0

Steps

- 1. Enter the Dell PERC 11 Configuration Utility. See Navigate to Dell PERC 11 configuration utility.
- 2. Click Main Menu > Configuration Management > Auto Configure RAID 0.
- **3.** Select **Confirm** and click **Yes** to continue. A RAID 0 Virtual disk is created on all physical disks that are in Ready state.

Create virtual disks

Steps

- 1. Enter the Dell PERC 11 Configuration Utility. See, Navigate to Dell PERC 11 configuration utility.
- 2. Click Main Menu > Configuration Management > Create Virtual Disk. The following list of options are displayed for you to define the virtual disk parameters:

Table 10. Create virtual disks

Option	Description
Create Virtual Disk	Allows you to create virtual disk selecting the RAID level, physical disks, and virtual disk parameters
Select RAID level	Allows you to choose the RAID level of your choice
Secure Virtual Disk	If you want to create a secured virtual disk, select Secure Virtual Disk. (i) NOTE: The Secure Virtual Disk option is enabled by default, only if the security key has been configured. Only SED physical disks are listed.
Select Physical Disks From	 Allows you to select one of the physical disk capacities: Unconfigured Capacity: creates a virtual disk on unconfigured physical disks. Free Capacity: utilizes unused physical disk capacity that is already part of a disk group.
Select Physical Disks	If you want to select the physical disks from which the virtual disks are being created, click Select Physical Disks . This option is displayed if you select Unconfigured Capacity as your physical disk capacity.
Select Disk Groups	If you want to select the disk groups from which the virtual disks are being created, click Select Disk Group . This option is displayed if you select Free Capacity as your physical disk capacity.
Configure Virtual Disk Parameters	Allows you to set the virtual disk parameters when creating the virtual disk. For more information, see Configuring virtual disk parameters.

3. Click Create Virtual Disk.

The virtual disk is created successfully.

NOTE: Ensure that you restart the system after creating a new Non-RAID or Virtual Disk on drives that previously had boot partitions.



Configure virtual disk parameters

Steps

- 1. Create a virtual disk, see Creating the virtual disks.
 - The Configure Virtual Disk Parameters section is displayed on the Create Virtual Disk screen.
- 2. In the Configure Virtual Disk Parameters section, you can set the following virtual disk parameters:

Table 11. Configure virtual disk parameters

Virtual disk parameters	Description
Virtual Disk Name	Allows you to enter the name for the virtual disk (i) NOTE: Allowed characters are A-Z, a-z, 0-9, underscore (_), and hyphen (-) only.
Virtual Disk Size	Displays the maximum capacity available for the virtual disk
Virtual Disk Size Unit	Displays the virtual disk storage space in megabytes, gigabytes, and terabyte.
Strip Element Size	Allows you to select the strip element size The disk striping involves partitioning each physical disk storage space in stripes of the following sizes: 64 KB, 128 KB, 256 KB, 512 KB, and 1 MB. By default, the strip element size is set to 256 KB.
Read Policy	 Displays the controller read policy You can set the read policy to: No read ahead—specifies that the controller does not use read ahead for the current virtual disk. Read ahead—specifies that the controller uses read ahead for the current virtual disk. Read ahead capability allows the controller to read sequentially ahead of requested data and store the additional data in the cache memory, anticipating that the data is required soon. By default, the read cache policy is set to read ahead.
Write Policy	 Displays the controller write cache policy You can set the write policy to: Write through—the controller sends a data transfer completion signal to the host when the disk subsystem has received all the data in a transaction. Write back—the controller sends a data transfer completion signal to the host when the controller cache has received all the data in a transaction. By default, the write policy is set to Write Back.
Disk Cache	Allows you to set the disk cache policy to default, enable, or disable. By default, the disk cache is set to default.
Default Initialization	 Displays the virtual disk initialization options. You can set the default initialization to: No — The virtual disk is not initialized. Fast — The first 8 MB of the virtual disk is initialized. Full — The entire virtual disk is initialized. For more information, see Virtual disk initialization. By default, the default initialization is set to No.

Create profile based virtual disk

Steps

- 1. Enter the Dell PERC 11 Configuration Utility. See Navigate to Dell PERC 11 configuration utility.
- 2. Click Main Menu > Configuration Management > Creating Profile Based Virtual Disk. The following list of RAID modes are displayed:
 - Generic RAID 0
 - Generic RAID 1
 - Generic RAID 5
 - Generic RAID 6
 - File Server



- Web/Generic Server
- Database
- **3.** Based on the RAID mode selected, one or more the physical disk selection criteria is displayed.
- **4.** From the **Physical Disk Selection Criteria** drop-down box, select a criterion based your requirement. The Profile Parameters of the selected option is displayed.
- 5. Click Create Virtual Disk
- Select Confirm and click Yes to continue. The virtual disk is created with the parameters of the profile selected.

View disk group properties

Steps

- 1. Enter the **Dell PERC 11 Configuration Utility**. See Navigate to Dell PERC 11 configuration utility.
- Click Main Menu > Configuration Management > View Disk Group Properties. The list of disk group properties are displayed:

Table 12. View disk group properties

Properties	Descriptions
Capacity Allocation	Displays all the virtual disks associated with the specific disk group. It also provides information about the available free space
Secured	Displays whether the disk group is secured or not

Convert to Non-RAID disk

Prerequisites

To convert a physical disk to non–RAID disk from the HII Configuration Utility, perform the following steps:

Steps

- 1. Enter the Dell PERC 11 Configuration Utility. See Navigate to Dell PERC 11 configuration utility.
- Click Main Menu > Configuration Management > Convert to Non-RAID Disk. The list of physical disks appears.
- 3. Select the physical disk to convert to Non-RAID disk.
- 4. Click Ok.
- A screen appears asking if you are sure you want to perform the operation.
- 5. Select the Confirm option.
- 6. Click Yes.

The operation is successful.

Delete configurations

Steps

- 1. Enter the Dell PERC 11 Configuration Utility. See Navigate to Dell PERC 11 configuration utility.
- 2. Click Main Menu > Configuration Management > Clear Configuration. A screen is displayed asking if you are sure you want to perform the operation.
- 3. CAUTION: It is recommended that you back up data stored on the virtual disks and hot spare disks on the controller before deleting the virtual drive.

Select **Confirm** and click **Yes** to continue. The virtual disks and hot spare disks available on the controller are deleted successfully.



Controller management

Clear controller events

Steps

- 1. Enter the Dell PERC 11 Configuration Utility. See Navigate to Dell PERC 11 configuration utility.
- 2. Click Main Menu > Controller Management > Advanced Controller Management.
- Click Clear Controller Events.
 A screen is displayed asking if you are sure you want to clear the controller events.
- 4. Select Confirm and click Yes to continue.

Save controller events

Steps

- 1. Enter the Dell PERC 11 Configuration Utility. See Navigate to Dell PERC 11 configuration utility.
- 2. Click Main Menu > Controller Management > Advanced Controller Management.
- Click Save Controller Events.
 A screen is displayed asking if you want to replace the existing file name.
- 4. Select Confirm and click Yes to continue.

Save debug log

Steps

- 1. Enter the Dell PERC 11 Configuration Utility. See Navigate to Dell PERC 11 configuration utility.
- 2. Click Main Menu > Controller Management > Advanced Controller Management.
- 3. Click Save Debug Log.
 - A screen is displayed indicating that the operation is successful.
- 4. Click **Ok**.

Enable security

Steps

- 1. Enter the Dell PERC 11 Configuration Utility. See Navigate to Dell PERC 11 configuration utility.
- 2. Click Main Menu > Controller Management > Advanced Controller Management.
- 3. Click Enable security, select Local Key Management.
- 4. Click Ok.
- 5. If you want to use the passphrase generated by the controller, click **Suggest Passphrase** and **Confirm** the passphrase by re-entering.
 - The operation is successful.
- Select I Recorded the Security Settings For Future Reference, click Enable Security. A screen is displayed indicating that the security will be enabled on this controller if you proceed.
- 7. Select **Confirm** and click **Yes** to continue. The operation is successful and click **Ok**.

Disable security

Steps

1. Enter the Dell PERC 11 Configuration Utility. See Navigate to Dell PERC 11 configuration utility.



- 2. Click Main Menu > Controller Management > Advanced Controller Management.
- 3. Click Disable security.
 - A screen is displayed asking if you are sure you want to disable security.
- Select Confirm and click Yes to continue. The operation is successful and click Ok.

Change security settings

Steps

- 1. Enter the Dell PERC 11 Configuration Utility. See Navigate to Dell PERC 11 configuration utility.
- 2. Click Main Menu > Controller Management > Advanced Controller Management.
- 3. Click Change Security Settings, select Change Current Security Settings.
- 4. Click Ok.
- 5. If you want to use the passphrase generated by the controller, click **Suggest Passphrase** and **Confirm** the passphrase by re-entering.
- The operation is successful.
- 6. Click Save Security Settings.
- Select Confirm and click Yes to continue. The operation is successful and click Ok.

Restore factory default settings

Steps

- 1. Enter the Dell PERC 11 Configuration Utility. See Navigate to Dell PERC 11 configuration utility.
- 2. Click Main Menu > Set Factory Defaults.
- A screen is displayed asking you to confirm the operation.
- 3. Select Confirm and click Yes to continue.

Auto configure behavior

Steps

- 1. Enter the Dell PERC 11 Configuration Utility. See Navigate to Dell PERC 11 configuration utility.
- 2. Click Main Menu > Controller Management > Advanced Controller Management > Manage Controller Mode. You can view the current Controller Mode.
- 3. Click Manage Controller Mode.
 - If required, you can view or change the hard drive settings for the controller. The possible options are:
 - Off and Non-RAID Disk
- 4. Click Apply Changes to save the changes.
- 5. Select Confirm and click Yes to continue.

NOTE: This feature is supported on PERC H355 adapter SAS, PERC H355 front SAS, PERC H350 Mini Monolithic SAS, and PERC H350 adapter SAS.

Manage controller profile

About this task

View the details of the profile and choose the desired profile, if supported. To view the properties of the controller profiles:

Steps

1. Enter the Dell PERC 11 Configuration Utility. See Navigate to Dell PERC 11 configuration utility.

Manage PERC 11 controllers using HII configuration utility 69



2. Click Main Menu > Controller Management > Advanced Controller Management > Manage Controller Profiles. The current profile and profile properties are displayed.

Advanced controller properties

Set the patrol read mode

Steps

- 1. Enter the Dell PERC 11 Configuration Utility. See Navigate to Dell PERC 11 configuration utility.
- 2. Click Main Menu > Controller Management > Advanced Controller Properties.
- 3. Click Patrol Read.
 - The following options are displayed:
 - Start—Starts patrol read for the selected controller.
 - Suspend—Suspends the ongoing patrol read operation on the controller.
 - Resume—Resumes the suspended patrol read operation.
 - Stop—Stops patrol read for the selected controller.
- 4. Set the Mode to Auto, Manual, or Disabled.
- 5. Click Apply Changes.

Enable physical disk power management

Steps

- 1. Enter the Dell PERC 11 Configuration Utility. See Navigate to Dell PERC 11 configuration utility.
- 2. Click Main Menu > Controller Management > Advanced Controller Properties.
- 3. Click Physical Disk Power Management.

The following list of options is displayed:

- Time Interval for Spin Down—allows the user to specify the delay time before a disk is spun down.
- Spin Down Hot Spare—allows you to enable or disable the spin down of hot spare disks.
- Spin Down Unconfigured Good—spin down of un-configured disks.
- Select the applicable options and click Apply Changes. The changes made are saved successfully.

Configure hot spare drives

Steps

- 1. Enter the Dell PERC 11 Configuration Utility. See Navigate to Dell PERC 11 configuration utility.
- 2. Click Main Menu > Controller Management > Advanced Controller Properties.
- 3. Click Spare.
 - The following list of options are displayed:
 - Persistent Hot Spare—allows you to enable or disable the ability to have same system backplane or storage enclosure disk slots dedicated as hot spare slots.
 - Allow Replace Member with Revertible Hot Spare—allows you to enable or disable the option to copy the data form a hot spare disk to physical disk.
 - Auto Replace Member on Predictive Failure—allows you to enable or disable the option to start a Replace Member operation if a predictive failure error is detected on a physical disk.
- 4. Select the applicable option and click **Apply Changes**.

The changes made are saved successfully.



Set task rates

Steps

- 1. Enter the Dell PERC 11 Configuration Utility. See Navigate to Dell PERC 11 configuration utility.
- 2. Click Main Menu > Controller Management > Advanced Controller Properties.
- 3. Click Task Rates.
 - The following options are displayed:
 - Background Initialization (BGI) Rate
 - Consistency Check Rate
 - Rebuild Rate
 - Reconstruction Rate
- 4. You can make the necessary changes and then click **Apply Changes**. The task rates operation is completely successfully.

Properties of Enterprise Key Management (EKM)

Steps

- 1. Enter the Dell PERC 11 Configuration Utility. See Navigate to Dell PERC 11 configuration utility.
- 2. Click Main Menu > Controller Management > Advanced Controller Properties.
- **3.** Click **Enterprise Key Management**. The properties of Enterprise Key Management is displayed.

Controller properties

Auto import foreign configuration

Steps

- 1. Enter the Dell PERC 11 Configuration Utility. See Navigate to Dell PERC 11 configuration utility.
- 2. Click Main Menu > Controller Management > Advanced Controller Properties.
- 3. In the Controller Properties section, set the Auto Import Foreign Configuration option to Enabled or Disabled.
- 4. Click Apply Changes.

Disable auto import

Steps

- 1. Enter the Dell PERC 11 Configuration Utility. See Navigate to Dell PERC 11 configuration utility.
- 2. Click Main Menu > Controller Management > Advanced Controller Properties.
- 3. In the Controller Properties section, set the Auto Import Foreign Configuration option to Disabled.
- **4.** Click **Apply Changes**. The auto import is disabled successfully.

Enable auto import

Steps

- 1. Enter the Dell PERC 11 Configuration Utility. See Navigate to Dell PERC 11 configuration utility.
- 2. Click Main Menu > Controller Management > Advanced Controller Properties.
- 3. In the Controller Properties section, set the Auto Import Foreign Configuration option to Enabled.
- 4. Click Apply Changes. The auto import is enabled successfully.



Select boot mode

Steps

- 1. Enter the Dell PERC 11 Configuration Utility. See Navigate to Dell PERC 11 configuration utility.
- 2. Click Main Menu > Controller Management > Advanced Controller Properties.
- **3.** In the **Controller Properties** section, select boot mode from the **Boot Mode** drop-down box. The following lists of boot mode options appear:

Table 13. Boot mode options

Option	Description
Stop on errors	The system stops during boot for errors which require attention from the user to rectify the issue.
Pause on errors	System pauses during boot to show errors but continue boot after it times out. Only critical events with an infinite timeout halt boot and require the user's attention to correct the issue.

NOTE: In UEFI BIOS mode, errors with timeouts do not appear during boot. It is designed to arise only in legacy BIOS mode.

(i) NOTE: By default, the boot mode option is set to pause on errors.

4. Click Apply Changes.

The boot mode operation is completed successfully.

Abort the consistency check

Steps

- 1. Enter the Dell PERC 11 Configuration Utility. See Navigate to Dell PERC 11 configuration utility.
- 2. Click Main Menu > Controller Management > Advanced Controller Properties.
- 3. In the Controller Properties section, set the Abort Consistency Check on Error option to Enabled or Disabled.
- 4. Click Apply Changes.
- The option to abort the consistency check operation on a redundant virtual disk is enabled if there is any inconsistency found in the data.

Preboot trace buffer

Steps

- 1. Enter the Dell PERC 11 Configuration Utility. See Navigate to Dell PERC 11 configuration utility.
- 2. Click Main Menu > Controller Management > Advanced Controller Properties.
- 3. In the Controller Properties section, set the Preboot Trace Buffer option to Enabled or Disabled.
- 4. Click Apply Changes.

Clear the cache memory

Steps

- 1. Enter the Dell PERC 11 Configuration Utility. See Navigate to Dell PERC 11 configuration utility.
- 2. Click Main Menu > Controller Management > Advanced Controller Properties.
- **3.** Click **Cache and Memory** > **Discard Preserved Cache**. The preserved cache is cleared successfully.
- 72 Manage PERC 11 controllers using HII configuration utility



Enable boot support

Steps

- 1. Enter the Dell PERC 11 Configuration Utility. See Navigate to Dell PERC 11 configuration utility.
- 2. Click Main Menu > Controller Management.
- 3. From the Select Boot Device drop-down box, select the primary bootable device.

In **Select Boot Device**, you will not be able to view 4 K sector drives. To view all the virtual disks created, navigate to the **Virtual Disk Management** screen in HII. For more information, see Virtual disk management.

If no boot device is selected, the first virtual disk will be set as the boot device on the next reboot. A Non-RAID disk is auto-selected as the boot device, if the controller does not have any virtual disks present.

- (i) NOTE: Select Boot Device is only applicable in legacy BIOS mode.
- (i) NOTE: 4 K sector drives boot support is only available in UEFI mode and managed by the boot loader.
- 4. Click Apply Changes.

Boot support is enabled for the selected controller.

Virtual disk management

Virtual disk numbering

Virtual disks are numbered in descending order beginning with the highest, ID 239.

View virtual disk properties

Steps

- 1. Enter the Dell PERC 11 Configuration Utility. See Navigate to Dell PERC 11 configuration utility.
- Click Main Menu > Virtual Disk Management.
 All the virtual disks associated with the RAID controller are displayed.
- **3.** To view the properties, click on the virtual disk. You can view the following properties of the Virtual disk:

Table 14. Virtual disk properties

Option	Description
Operation	List of operations you can perform on the selected virtual disk. The options are: Blink Unblink Delete Virtual Disk Reconfigure Virtual Disks Fast Initialization Slow Initialization
Name	Indicates the name of the virtual disk.
RAID level	Indicates the RAID level of the virtual disk.
Status	Indicates the status of the virtual disk. The possible options are: Optimal Degraded Offline Failed
Size	Indicates the size of the virtual disk.



View physical disks associated with a virtual disk

Steps

- 1. Enter the Dell PERC 11 Configuration Utility. See Navigate to Dell PERC 11 configuration utility.
- 2. Click Main Menu > Virtual Disk Management. All the virtual disks associated with the RAID controller are displayed.
- **3.** Click on a virtual disk. The properties of the virtual disk are displayed.
- Click View Associated Physical Disks. All the physical disks that are associated with the virtual disk are displayed.
- 5. From the **Associated Physical Disks** section, select the physical disk.
- 6. Click View Physical Disk Properties to view the physical disk properties.

View advanced properties of a virtual disk

Steps

- 1. Enter the Dell PERC 11 Configuration Utility. See Navigate to Dell PERC 11 configuration utility.
- 2. Click Main Menu > Virtual Disk Management.

All the virtual disks associated with the RAID controller are displayed.

- **3.** Click the virtual disk. The properties of the virtual disk are displayed.
- 4. Click Advanced....
 - You can view the following additional properties of the virtual disk:

Table 15. Advanced properties of the virtual disk

Option	Description
Logical sector size	Indicates the logical sector size of this virtual disk.
Strip element size	Indicates the strip element size for the virtual disk.
Secured	Indicates whether the virtual disk is secured or not.
Bad blocks	Indicates whether the virtual disk has corrupted blocks.

Configure virtual disk policies

Steps

- 1. Enter the Dell PERC 11 Configuration Utility. See Navigate to Dell PERC 11 configuration utility.
- 2. Click Main Menu > Virtual Disk Management.
 - All the virtual disks associated with the RAID controller are displayed.
- 3. Click Advanced....
 - You can view the following virtual disk policies:

Table 16. Virtual disk policies

Option	Description
Current write cache	Indicates the current write cache policy for the virtual disk.
Default write cache	 Allows selection of the write cache policy for the virtual disk. The possible options are: Write Through Write Back Force Write Back
Read cache policy	Allows selection of the read cache policy for the virtual disk. The possible options are:

74 Manage PERC 11 controllers using HII configuration utility



Table 16. Virtual disk policies (continued)

Option	Description
	No Read AheadRead Ahead
Disk cache	 Allows selection of the disk cache policy for the virtual disk. The possible options are: Default (Disk Default) Enable Disable

4. Click Apply Changes.

The changes made are saved successfully.

Configure Virtual Disks

When configuring the virtual disks, you should consider the workload intended; RAID 1: for simple boot disk; RAID 5 or 6: for file or web servers (sequential reads/writes of files); RAID 10: for transactional database (small random reads and writes).

Virtual disks configured on hard drives should use the controller default cache setting of Write Back and Read Ahead.

Virtual disks configured on SSDs can use the same controller defaults settings as hard drives. Most users perform a copy of OS files or a data base to the new array. This setting provides optimum performance in this configuration.

Once the copy is complete, the array can be used as it is depending on the number and type of SSDs. It is recommended to enable FastPath by changing the controller's Write cache policy to Write Through and the Read cache policy to No Read Ahead. FastPath is developed to achieve the best random read/write performance from SSDs.

Only IO block sizes smaller than the virtual disk's stripe size are eligible for FastPath. In addition, there should be no background operations (rebuild, initialization) running on the virtual disks. FastPath is disabled if there is active background operation.

(i) NOTE: RAID 50, and RAID 60 virtual disks cannot use FastPath.

(i) NOTE: The Physical Disk Power Management feature is not applicable to FastPath-capable virtual disks.

Perform expand virtual disk operation

Prerequisites

To enable expand virtual disk feature from the HII Configuration Utility, perform the following steps:

Steps

- 1. Enter the Dell PERC 11 Configuration Utility. See Navigate to Dell PERC 11 configuration utility.
- 2. Click Main Menu > Virtual Disk Management.
- The list of virtual disks is displayed.
- **3.** Select the virtual disk.
- 4. From the Operations drop-down menu, select Expand Virtual Disk.

(i) NOTE: You can view the Expand Virtual Disk feature only if there is free space available in the associated disk group.

- 5. Click Go.
- 6. To expand virtual disk, enter the percentage of available capacity, and then click **Ok**. A screen is displayed asking if you are sure you want to perform the operation.
- 7. Select the **Confirm** option.
- 8. Click Yes.

The expand virtual disk operation is completed successfully.



Perform consistency check

Prerequisites

To enable consistency check from the HII Configuration Utility, perform the following steps:

Steps

- 1. Enter the Dell PERC 11 Configuration Utility. See Navigate to Dell PERC 11 configuration utility.
- 2. Click Main Menu > Virtual Disk Management.
- The list of virtual disks is displayed.
- **3.** Select the virtual disk.

(i) NOTE: Consistency check cannot be run on RAID 0 virtual disks.

- 4. From the Operations drop-down menu, select Check Consistency.
- Click Go.
 A screen is displayed asking if you are sure you want to perform the operation.
- 6. Select the **Confirm** option.
- 7. Click Yes.
 - The consistency check operation is completed successfully.

Physical disk management

View physical disk properties

Steps

- 1. Enter the Dell PERC 11 Configuration Utility. See Navigate to Dell PERC 11 configuration utility.
- Click Main Menu > Physical Disk Management. All the physical disks that are associated with the RAID controller are displayed.
- 3. To view the properties, click the physical disk.

Table 17. Physical disk properties

Option	Description
Operation	 The list of operations you can perform on the selected physical disk. The options are: Blink Unblink Assign global hot spare Cryptographic erase Convert to non-RAID disk
Device ID	Unique identifier of the physical disk.
Backplane ID	Backplane ID in which the physical disk is located in for PERC H755 adapter, PERC H755 front SAS, PERC H755N front NVMe, PERC H750 adapter SAS, PERC H755 MX adapter, PERC H355 adapter SAS, PERC H355 front SAS, PERC H350 adapter SAS , and PERC H350 Mini Monolithic SAS
Slot number	The drive bay in which the physical disk is located for the corresponding backplane or enclosure to which the controller is connected.
Status	Status of the physical disk.
Size	Size of the physical disk.
Туре	Type of the physical disk.
Model	Model of the physical disk.

76 Manage PERC 11 controllers using HII configuration utility



Table 17. Physical disk properties (continued)

Option	Description
Serial number	Serial of the physical disk.

4. To view additional properties of the physical disk, click Advanced....

Table 18. Advanced physical disk properties

Option	Description
Logical sector size	Logical sector size of the selected physical disk
Physical sector size	Physical sector size of the selected physical disk
SMART status	SMART status of a physical disk
Revision	Firmware version of the physical disk
WWID	Unique identifier used to identify the device
Multipath	Multipath of the controller
Physical disk power state	Power condition (On or Power Save) of the physical disk
Disk cache setting	Disk cache setting (i) NOTE: Disk cache for SATA Gen3 drives is disabled by default.
Disk protocol	Type of hard disk used
Device speed	Speed of the physical disk
Negotiated link speed	Negotiated link speed of the device
PCIe capable link width	N/A for SAS/SATA drives
PCIe negotiated link width	N/A for SAS/SATA drives
Encryption capable	Encryption capability of the physical disk
Encryption supported	Encryption capability enabled at the controller level
Secured	Security status of the physical disk
Cryptographic erase capable	Cryptographic erase capability of the physical disk

Cryptographic erase

Cryptographic erase is a process to erase all data permanently on an encryption-capable and unconfigured physical disk, and reset the security attributes.

Prerequisites

- The non-RAID and virtual disks associated with the drive are deleted.
- The disks are not hot spares.

About this task

The Cryptographic erase feature is supported only on Instant Secure Erase (ISE) and Self Encrypting Drives (SED) drives.

Steps

- 1. Enter the **Dell PERC 11 Configuration Utility**. See Navigate to Dell PERC 11 configuration utility.
- 2. Click Main Menu > Physical Disk Management.
- The list of physical disks is displayed.
- **3.** Select a physical disk.
- 4. From the Operations drop-down menu, select Cryptographic Erase.



(i) NOTE: If the drive installed is ISE or SED capable only then the Cryptographic erase option is displayed.

5. Click Go.

A screen is displayed asking if you are sure you want to perform the operation.

- 6. Select the **Confirm** option.
- 7. Click Yes.

The Cryptographic erase operation is completed successfully.

Physical disk erase

Prerequisites

To use the Physical Disk Erase feature from the HII Configuration Utility, perform the following steps:

Steps

- 1. Enter the **Dell PERC 11 Configuration Utility**. See Navigate to Dell PERC 11 configuration utility.
- Click Main Menu > Physical Disk Management. The list of physical disks is displayed.
- **3.** Select a physical disk.
- 4. From the Operations drop-down menu, select Physical Disk Erase.

(i) NOTE: If the drive installed is neither SED or ISE capable, then only the Physical Disk Erase option is displayed.

5. Click Go.

A screen is displayed asking if you are sure you want to perform the operation.

- 6. Select the Confirm option.
- 7. Click Yes.

The physical disk erase operation is completed successfully.

Assigning a global hot spare

Prerequisites

To assign a global hot spare from the HII Configuration Utility, perform the following steps:

Steps

- 1. Enter the Dell PERC 11 Configuration Utility. See Navigate to Dell PERC 11 configuration utility.
- 2. Click Main Menu > Physical Disk Management.
- The list of physical disks is displayed.
- **3.** Select the physical disk.
- 4. From the Operations drop-down menu, select Assign Global Hot Spare.
- 5. Click Go.
 - A screen is displayed asking if you are sure you want to perform the operation.
- 6. Select the Confirm option.
- 7. Click Yes.

The global hot spare disk is created successfully.

Assigning a dedicated hot spare

Prerequisites

To assign a dedicated hot spare from the HII Configuration Utility, perform the following steps:



Steps

- 1. Enter the Dell PERC 11 Configuration Utility. See Navigate to Dell PERC 11 configuration utility.
- 2. Click Main Menu > Physical Disk Management.
- The list of physical disks is displayed.
- **3.** Select the physical disk.
- 4. From the Operations drop-down menu, select Assign Dedicated Hot Spare.
- 5. Click Go.
- A screen is displayed asking if you are sure you want to perform the operation.
- 6. Select the **Confirm** option.
- 7. Click Yes.

The dedicated hot spare disk is created successfully.

Convert to RAID capable

Prerequisites

To convert a non-RAID disk to RAID capable disk from the HII Configuration Utility, perform the following steps:

Steps

- 1. Enter the Dell PERC 11 Configuration Utility. See Navigate to Dell PERC 11 configuration utility.
- 2. Click Main Menu > Physical Disk Management. The list of physical disks appears.
- **3.** Select the physical disk.
- 4. From the Operations drop-down menu, select Convert to RAID capable.
- 5. Click Go.
- A screen appears asking if you are sure you want to perform the operation.
- $\textbf{6.} \hspace{0.1 cm} \text{Select the } \textbf{Confirm} \hspace{0.1 cm} \text{option}.$
- 7. Click Yes. The operation is successful.

Convert to Non-RAID disk

Prerequisites

To convert a physical disk to non-RAID disk from the HII Configuration Utility, perform the following steps:

Steps

- 1. Enter the Dell PERC 11 Configuration Utility. See Navigate to Dell PERC 11 configuration utility.
- 2. Click Main Menu > Physical Disk Management. The list of physical disks appears.
- 3. Select the physical disk.
- 4. From the Operations drop-down menu, select Convert to Non-Raid disk.
- 5. Click Go.
 - A screen appears asking if you are sure you want to perform the operation.
- 6. Select the Confirm option.
- 7. Click Yes. The operation is successful.



Hardware components

View battery properties

Steps

- 1. Enter the Dell PERC 11 Configuration Utility. See Navigate to Dell PERC 11 configuration utility.
- 2. Click Main Menu > Hardware Components > Battery Management.
- The battery and capacity information are displayed.
- $\ensuremath{\textbf{3}}.$ You can view the following properties of the battery:

Table 19. Battery properties

Field	Description
Туре	Displays the type of battery available.
Status	Displays the current status of the battery.
Temperature	Displays the current temperature of the battery and also indicates whether the temperature is normal or high.
Charge	Displays the available charge of the battery in percentage.

4. Displays click Advanced....

The additional advanced properties of the physical battery are displayed.

5. You can view the following advanced properties of the battery:

Table 20. Advanced battery properties

Field	Description
Status	Displays whether the current status of the battery is learning, degraded, or failed.
Voltage	Displays whether the voltage status of the battery is normal or high.
Current	Displays power consumption of the battery in milliamps (mA).
Full capacity	Displays the maximum charge capacity of the battery.
Remaining capacity	Displays the current charge capacity of the battery.
Expected margin of error	Displays expected margin of error.
Completed discharge cycles	Displays the completed discharge cycles.
Learn mode	Displays the condition of the battery. The learn cycle is a periodic operation that calculates the charge that is remaining in the battery to ensure there is sufficient energy.

View physical disks associated with an enclosure

Steps

- 1. Enter the Dell PERC 11 Configuration Utility. See Navigate to Dell PERC 11 configuration utility.
- 2. Click Main Menu > Hardware Components > Enclosure Management.
- 3. From the Select Enclosure field, choose the enclosure for which you need to view the physical disks.
- All the physical disks that are associated with the virtual disk are displayed.
- Click the Attached Physical Disks drop-down box.
 All the physical disks that are associated with the selected enclosure are displayed.
- 80 Manage PERC 11 controllers using HII configuration utility



Security key management in HII configuration utility

The Dell OpenManage storage management application and the **HII Configuration Utility** of the controller allow security keys to be created and managed as well as create secured virtual disks. The following section describes the menu options specific to security key management and provide detailed instructions to perform the configuration tasks. The contents in the following section apply to the **HII Configuration Utility**. For more information on the management applications, see Applications and User Interfaces supported by PERC 11.

- The **Controller Management** screen displays controller information and action menus. You can perform the following security-related actions through the controller management menu:
 - **Security Key Management**—Creates or changes the local key management (LKM) security key. Deletes the local key management (LKM) or secure enterprise key manager (SEKM) security key.
- The **Virtual Disk Management** screen displays physical disk information and action menus. You can perform the following security related actions through the virtual disk management menu:
 - Secure Disk Group—Secures all virtual disks in disk group.
 - **Create secure virtual disk**—Creates a new virtual disk that is secured with the security key on the controller.
- The **Physical Disk Management** screen displays physical disk information and action menus. You can perform the following security-related actions through the physical disk management menu:
 - Secure non-RAID disk—Secures the non-RAID disk with the controller security key.
 - **Cryptographic Erase**—Permanently erases all data on the physical disk and resets the security attributes.

For more information on the Physical Disk Management screen and the Virtual Disk Management screen, see Physical disk management and Virtual disk management.



Security key and RAID management

Topics:

- Security key implementation
- Local Key Management
- Create a security key
- Change Security Settings
- Disable security key
- Create a secured virtual disk
- Secure a non-RAID disk
- Secure a pre-existing virtual disk
- Import a secured non-RAID disk
- Import a secured virtual disk
- Dell Technologies OpenManage Secure Enterprise Key Manager

Security key implementation

The PERC 11 series of cards support self-encrypting disk (SED) for protection of data against loss or theft of SEDs. Protection is achieved by the use of encryption technology on the drives. There is one security key per controller. You can manage the security key using local key management (LKM) or OpenManage Secure Enterprise Key Manager, also referred as Secure Enterprise Key Manager (SEKM). The LKM key can be escrowed in to a file using Dell OpenManage Storage Management application. The security key is used by the controller to lock and unlock access to encryption-capable physical disks. In order to take advantage of this feature, you must:

- **1.** Have SEDs in your system.
- **2.** Create a security key.
- **NOTE:** If the host system is powered off when connected to an external enclosures or if the sled is powered off in C6XXX PowerEdge servers, the drives will remain in an unlocked state until they are power cycled or AC power is disconnected from the sled or external enclosure.

Local Key Management

You can use Local Key Management (LKM) to generate the key ID and the passphrase that is required to secure the virtual disk. You can secure virtual disks, change security keys, and manage secured foreign configurations using this security mode. **NOTE:** LKM mode is not supported on PERC H355 adapter SAS, PERC H350 adapter SAS, PERC H355 front SAS, and PERC H350 Mini Monolithic SAS.

Create a security key

About this task

(i) NOTE: There is no passphrase backup option when you create a security key; you need to remember your passphrase.

Steps

- 1. Enter the Dell PERC 11 Configuration Utility. See Navigate to Dell PERC 11 configuration utility.
- 2. Click Main Menu > Controller Management > Advanced Controller Management > Enable Security.
- 3. Select the Security Key Management mode as Local Key Management.

Inserido ao protocolo **22.951.206-4** por: **Pedro Henrique Golin Linhares** em: 31/03/2025 12:49. A autenticidade deste documento pode ser validada no endereço: https://www.eprotocolo.pr.gov.br/spiweb/validarDocumento com o código: **692a5ef82beb173ed33dad4ee6505546**.



4. Click Ok.

5. In the Security Key Identifier field, enter an identifier for your security key.

NOTE: The Security Key Identifier is a user supplied clear text label used to associate the correct security key with the controller.

- 6. If you want to use the passphrase generated by the controller, click **Suggest Passphrase**. Assigns a passphrase suggested by the controller automatically.
- 7. In the **Passphrase** field, enter the passphrase.

NOTE: Passphrase is case-sensitive. You must enter minimum 8 or maximum 32 characters. Ensure that the characters contain at least one number, one lower case letter, one upper case letter, and one non-alphanumeric character.

8. In the **Confirm** field, re-enter the passphrase to confirm.

NOTE: If the Passphrase entered in the Passphrase and Confirm fields do not match, then you are prompted with an error message to enter the passphrase again.

- 9. Select the I recorded the Security Settings for Future Reference option.
- 10. Click Enable Security.

The Security Key is created successfully.

Change Security Settings

Steps

- 1. Enter the Dell PERC 11 Configuration Utility. See Navigate to Dell PERC 11 configuration utility.
- 2. Click Main Menu > Controller Management > Advanced Controller Management > Change Security Settings.
- 3. Select security identifier:
 - a. To change the Security key Identifier enter a new key identifier in Enter a New Security Key identifier text box.
 - b. To keep existing key identifier, select Use the existing Security Key Identifier check box.
- 4. Enter the existing passphrase.
- 5. Set passphrase:
 - a. To change the security passphrase, enter a new passphrase in the **Enter a New Passphrase** text box. Re-enter the new passphrase to confirm.
 - b. To keep the existing passphrase, select Use the existing passphrase.
- 6. Select I recorded the Security Settings for Future Reference.
- 7. Click Save Security Settings.
- 8. Select **Confirm** and then click **Yes**. Security settings changed successfully.

Disable security key

About this task

(i) NOTE: Disabling Security Key is active if there is a security key present on the controller.

Steps

- 1. Enter the Dell PERC 11 Configuration Utility. See Navigate to Dell PERC 11 configuration utility.
- 2. Click Main Menu > Controller Management > Advanced Controller Management > Disable Security. You are prompted to confirm whether you want to continue.
- 3. Select the Confirm option.
- 4. Click Yes.
 - The security key is disabled successfully.

(i) NOTE: All virtual disks must be deleted or removed to disable security.



WARNING: Any un-configured secured disks in the system will be repurposed.

Create a secured virtual disk

About this task

To create a secured virtual disk, the controller must have a security key established first. See Create a security key.

NOTE: Combining SAS and SATA hard drives within a virtual disk is not supported. Also, combining hard drives and olid-state drives (SSDs) within a virtual disk is not supported. Mixing of NVMe drives is not supported.

After the security key is established, perform the following steps:

Steps

- 1. Enter the Dell PERC 11 Configuration Utility. See Navigate to Dell PERC 11 configuration utility.
- Click Main Menu > Configuration Management > Create Virtual Disk. For more information, see Create virtual disks.
- 3. Select the Secure Virtual Disk option.
- Click Create Virtual Disk. The secure virtual disk is created successfully.

Secure a non-RAID disk

In HII, secure a non-RAID disk by using the security key of the controller.

Steps

- 1. Enter the Dell PERC 11 Configuration Utility. See Navigate to Dell PERC 11 configuration utility.
- 2. Click Main Menu > Physical Disk Management.
- The list of Non-RAID disks is displayed.
- **3.** Select a non-RAID disk.
- 4. From the Operations drop-down menu, select Secure Non-RAID Disk.

Secure a pre-existing virtual disk

Steps

- 1. Enter the Dell PERC 11 Configuration Utility. See Navigate to Dell PERC 11 configuration utility.
- Click Main Menu > Virtual Disk Management. The list of virtual disks is displayed.
- 3. Select a virtual disk.
- 4. From the Operations drop-down menu, select Secure Virtual Disk.

(i) NOTE: The virtual disks can be secured only when the virtual disks are in Optimal state.

Import a secured non-RAID disk

If you are inserting a non-RAID disk into a system that has a controller key different from the security key on the drive, the security key from the system in which it was initially secured must be provided in HII.

Prerequisites

(i) NOTE: The controller must have an existing security key before importing a secured non-RAID disk.

84 Security key and RAID management

Inserido ao protocolo **22.951.206-4** por: **Pedro Henrique Golin Linhares** em: 31/03/2025 12:49. A autenticidade deste documento pode ser validada no endereço: https://www.eprotocolo.pr.gov.br/spiweb/validarDocumento com o código: **692a5ef82beb173ed33dad4ee6505546**.



Steps

- 1. Enter the Dell PERC 11 Configuration Utility. See Navigate to Dell PERC 11 configuration utility.
- 2. Click Main Menu > Configuration Management > Manage Foreign Configurations.
- 3. Click Enter Passphrase for Locked Disks.
- A screen is displayed asking if you are sure you want to perform the operation.
- 4. Enter Passphrase if importing non-RAID disk with a different passphrase.
- 5. Select the **Confirm** option.
- 6. Click Yes.

(i) NOTE: If Auto-Configure for non-RAID Disks is enabled, the disk becomes a non-RAID disk. Else, it is unconfigured.

Import a secured virtual disk

Prerequisites

(i) NOTE: The controller must have an existing security key before importing secured foreign virtual disk.

Steps

- 1. Enter the Dell PERC 11 Configuration Utility. See Navigate to Dell PERC 11 configuration utility.
- 2. Click Main Menu > Configuration Management > Manage Foreign Configurations > Preview Foreign Configurations.
- 3. Click Import Foreign Configuration.
- A screen is displayed asking if you are sure you want to perform the operation.
- 4. Enter **Passphrase** if importing virtual disk with a different passphrase.
- 5. Select the Confirm option.
- 6. Click Yes.

The foreign configuration is imported successfully.

Dell Technologies OpenManage Secure Enterprise Key Manager

This feature allows the PERC to receive a security key from a remote server instead of saving the key on a local controller. This protects data on secured disks under the PERC if the disks or entire system is stolen. Refer to the www.dell.com/idracmanuals for more information on configuring OpenManage Secure Enterprise Key Manager, as well as Secure Sockets Layer (SSL) or Transport Layer Security (TLS) related configuration.

NOTE: Downgrade of PERC firmware to a firmware that does not support enterprise key management while enterprise key manager mode is enabled, is blocked.

- **NOTE:** When replacing a controller enabled with enterprise key management, lifecycle controller part replacement will re-configure the new controller to match the existing controller's configuration.
- **NOTE:** If key exchange fails during boot, view and correct any connection issues with the key server identified in the iDRAC lifecycle log. Then the system can be cold booted.

Supported controllers for OpenManage Secure Enterprise Key Manager

Enterprise key manager mode is supported on the PERC H755 adapter, PERC H755 front SAS, and PERC H755N front NVMe, and allows the creation of secured virtual disks and non–RAID disks. For more information about supported platforms, see Support Site.



Enterprise key manager mode is not supported on the PERC H755 MX adapter, PERC H355 front SAS, PERC H355 adapter SAS, PERC H350 adapter SAS, and PERC H350 Mini Monolithic SAS.

Manage enterprise key manager mode

iDRAC manages Enterprise key manager features. For instructions on enabling enterprise key manager mode, vidi **dell.com/** idracmanuals.

- () NOTE: If preserved cache is present, the controller does not allow OpenManage Secure Enterprise Key Manager (SEKM) mode to be enabled.
- **NOTE:** When enterprise key manager mode is enabled, the controller waits up to two minutes for iDRAC to send keys, after which the PERC continues to boot.
- **NOTE:** Transitioning a controller from Local Key Management (LKM) mode to SEKM mode is supported on firmware starting with version 52.16.1-4074.
- **NOTE:** iDRAC performs rotation of keys. Any attempt to rekey the controller through a different management application is not supported.

Disable enterprise key manager mode

Enterprise key manager mode can be disabled from any supported Applications & User Interfaces supported by PERC 11. For more information, see the management application's user's guide or see Disable security key.

Manage virtual disks in enterprise key manager mode

Virtual disks are managed in the same way in enterprise key manager mode as in local key manager mode. SED capable virtual disks can be secured during or after creation. See Create a secured virtual disk.

Manage non-RAID disks in enterprise key manager mode

Non–RAID disks are managed in the same way in enterprise key manager mode as in local key manager mode. SED capable non–RAID disks can be secured after creation. See Create a secured virtual disk.

Transition of drives from local key management to enterprise key management (without supported firmware for PERC and iDRAC)

Local key management drives can be transitioned to an enterprise key management enabled system, but the controller cannot be transitioned from local key management mode to enterprise key manager mode or the reverse without first disabling security on the controller. Perform the following steps to transition from local key management drives to enterprise key management:

Steps

- 1. Save the current local key management security key.
- 2. Shut down both systems.
- 3. Remove the local key management drives and reinsert them to the enterprise key manager enabled system.
- 4. Power on the enterprise key manager system.
- 5. Go to HII foreign configuration.
- 6. Enter the local key management keys for those drives.
- 7. Import the configuration.
 - **NOTE:** Once local key management drives are migrated to enterprise key manager, they cannot be migrated back to local key management mode. The drives have to be cryptographically erased to disable security and then converted back to local key management disks. For more information about performing this action, contact Support Site.
- 86 Security key and RAID management

Inserido ao protocolo **22.951.206-4** por: **Pedro Henrique Golin Linhares** em: 31/03/2025 12:49. A autenticidade deste documento pode ser validada no endereço: https://www.eprotocolo.pr.gov.br/spiweb/validarDocumento com o código: 692a5ef82beb173ed33dad4ee6505546.



Migrate of drives from local key management to enterprise key management (with supported firmware for PERC and iDRAC)

PERC enables transition from Local Key Management (LKM) mode to Secure Enterprise Key Manager (SEKM) mode without disabling LKM security first. For instructions on transitioning from LKM mode to SEKM mode, see iDRAC Manuals.

(i) NOTE: This feature is supported on firmware starting with version 51.16.0-4076.

The transition from LKM to SEKM on the controller fails if the following are true at time of attempt:

- Snapdump is present on PERC.
- Preserved cache is present on PERC.
- RAID level migration is in progress on PERC,
- Online capacity expansion is in progress on PERC.
- Sanitize on a physical disk is in progress.
- LKM key that does not match with the current key of PERC.
- PERC firmware does not support transition.

Security key and RAID management 87



Troubleshooting issues in PERC11 cards

To get help for resolving issues in your PERC11 series cards, you can contact your Dell Technical Service representative. **Topics:**

- Single virtual disk performance or latency in hypervisor configurations
- Configured disks removed or not accessible error message
- Dirty cache data error message
- Discovery error message
- Drive Configuration Changes Error Message
- Windows operating system installation errors
- Firmware fault state error message
- Foreign configuration found error message
- Foreign configuration not found in HII
- Degraded state of virtual disks
- Memory errors
- Preserved Cache State
- Security key errors
- General issues
- Physical disk issues
- SMART errors
- Replace member errors
- Linux operating system errors
- Drive indicator codes
- HII error messages
- System reports more drive slots than what is available
- World Wide Number on drive sticker is not the same in applications
- · Backplane firmware revision not changing in PERC interfaces after an update

Single virtual disk performance or latency in hypervisor configurations

Multi-initiator or hypervisor configurations running multiple I/O workloads to a single raid array may experience degraded performance or latency. This is caused by upper layers sending separate I/O workloads for each virtual machine to the storage subsystem which ends up being a random I/O workload to the under lying RAID array. For I/O workload configurations that require lower latency restrictions and higher I/O performance it may be beneficial to run fewer I/O workloads to individual RAID arrays or to use separate RAID arrays and physical disks for each I/O workload. Other considerations are making sure write-back, read ahead cache is enabled for rotational disks or using solid state drives (SSDs) to improve random I/O workload performance.

Performance degradation may also be observed when background operations such as initialization, consistency check, or reconstructions are running on the virtual disk. See your hypervisor storage best practices or performance best practices guides for additional configuration support.

Inserido ao protocolo **22.951.206-4** por: **Pedro Henrique Golin Linhares** em: 31/03/2025 12:49. A autenticidade deste documento pode ser validada no endereço: https://www.eprotocolo.pr.gov.br/spiweb/validarDocumento com o código: 692a5ef82beb173ed33dad4ee6505546.



Configured disks removed or not accessible error message

Error Message:	Some configured disks have been removed from your system or are no longer accessible. Check your cables and ensure all disks are present. Press any key or 'C' to continue.
Probable Cause:	The message indicates that some configured disks were removed. If the disks were not removed, they are no longer accessible. The cables from the PERC controller to the backplane might be improperly connected.
Corrective Action:	Check the cable connections and fix issues if any. Restart the system. If there are no cable problems, press any key or $<$ C $>$ to continue.

Dirty cache data error message

Error Message:	The following virtual disks are missing: (x). If you proceed (or load the configuration utility), these virtual disks will be removed from your configuration. If you wish to use them at a later time, they will have to be imported. If you believe these virtual disks should be present, please power off your system and check your cables to ensure all disks are present. Press any key to continue, or 'C' to load the configuration utility. The cache contains dirty data, but some virtual disks are missing or will go offline, so the cached data cannot be written to disk. If this is an unexpected error, then please power off your system and check your cables to ensure all disks are present. If you continue, the data in cache will be permanently discarded. Press 'X' to acknowledge and permanently destroy the cached data.
Probable Cause:	The controller preserves the dirty cache from a virtual disk if the disk becomes offline or is deleted because of missing physical disks. This message indicates that some configured disks were removed. If the disks were not removed, they are no longer accessible. The cables from the PERC controller to the backplane might be improperly connected.

CorrectiveCheck the cable connections and fix any problems. Restart the system. Use the HII configuration utility to
import the virtual disk or discard the preserved cache. For the steps to discard the preserved cache, see
Clear the cache memory.

Discovery error message

Error Message: A discovery error has occurred, please power cycle the system and all the enclosures attached to this system.
 Probable Cause: This message indicates that discovery did not complete within 120 seconds. The cables from the PERC controller to the backplane might be improperly connected.
 Corrective Check the cable connections and fix any problems. Restart the system.

Drive Configuration Changes Error Message

Error Message: Entering the configuration utility in this state will result in drive configuration changes. Press 'Y' to continue loading the configuration utility or please power off your system and check your cables to ensure all disks are present and reboot.



Probable Cause:	The message is displayed after another HII warning indicating there are problems with previously configured disks and you have chosen to accept any changes and continue. The cables from the PERC controller to the backplane might be improperly connected.
Corrective	Check the cable connections and fix any problems before restarting the system. If there are no cable

Action:

Check the cable connections and fix any problems before restarting the system. If there are no cable problems, press any key or <Y> to continue.

Windows operating system installation errors

Ensure that you perform the following step before installing Windows on 4 KB sector drives:

1. Read and understand the updates to the version of Windows that you have installed. You can find this information in the Microsoft help. For more information, see Microsoft support policy for 4 K sector hard drives in Windows.

Firmware fault state error message

Error Message:Firmware is in Fault State.Corrective
Action:Contact Global Technical Support.

Foreign configuration found error message

Error Message:	Foreign configuration(s) found on adapter. Press any key to continue, or 'C' to load the configuration utility or 'F' to import foreign configuration(s) and continue.
Probable Cause:	When a controller firmware detects a physical disk with existing foreign metadata, it flags the physical disk as foreign and generates an alert indicating that a foreign disk was detected.
Corrective Action:	Press <f></f> at this prompt to import the configuration (if all member disks of the virtual disk are present) without loading the HII Configuration Utility . Or press <c></c> to enter the HII Configuration Utility and either import or clear the foreign configuration.

Foreign configuration not found in HII

Error Message:	The foreign configuration message is present during POST but no foreign configurations are present in the foreign view page in HII configuration utility. All virtual disks are in an optimal state.
Corrective	Ensure all your PDs are present and all VDs are in optimal state. Clear the foreign configuration using HII configuration utility or Dell OpenManage Server Administrator Storage Management.
Action:	CAUTION: The physical disk goes to Ready state when you clear the foreign configuration.

If you insert a physical disk that was previously a member of a virtual disk in the system, and that disk's previous location has been taken by a replacement disk through a rebuild, you must manually remove the foreign configuration flag of the newly inserted disk.

Degraded state of virtual disks

A redundant virtual disk is in a degraded state when one or more physical disks have failed or are inaccessible. For example, if a RAID 1 virtual disk consists of two physical disks and one of them fails or become inaccessible, the virtual disk become degraded.

To recover a virtual disk from a degraded state, you must replace the failed physical disk and rebuild it. Once the rebuilding process is complete, the virtual disk state changes from degraded to optimal.

90 Troubleshooting issues in PERC11 cards

Inserido ao protocolo **22.951.206-4** por: **Pedro Henrique Golin Linhares** em: 31/03/2025 12:49. A autenticidade deste documento pode ser validada no endereço: https://www.eprotocolo.pr.gov.br/spiweb/validarDocumento com o código: 692a5ef82beb173ed33dad4ee6505546.



Memory errors

Memory errors can corrupt cached data, so the controllers are designed to detect and attempt to recover from the memory errors. Single-bit memory errors can be handled by the controller and do not disrupt normal operation. A notification is sent if the number of single-bit errors exceeds a threshold value.

Multi-bit errors are more serious as they result in corrupted data and data loss. The following are the actions that occur in the case of multi-bit errors:

- If a multi-bit error occurs while accessing data in the cache when the controller is started with dirty cache, the controller discards the cache contents. The controller generates a warning message to the system console to indicate that the cache was discarded and generates an event.
- If a multi-bit error occurs at run-time either in code/data or in the cache, the controller stops.
- The controller logs an event to the controller's internal event log and a message during POST is displayed indicating a multi-bit error has occurred.

(i) NOTE: In case of a multi-bit error, contact Global Technical Support.

Preserved Cache State

The controller preserves the dirty cache from a virtual disk if the virtual disk goes offline or is deleted because of missing physical disks. This preserved dirty cache is called **pinned cache** and is preserved until you import the virtual disk or discard the cache.

- 1. Import the virtual disk—Power off the system, re-insert the virtual disk and restore the system power. Use the **HII Configuration Utility** to import the foreign configuration.
- 2. Discard the preserved cache—See Clear the cache memory.
- **NOTE:** It is recommended to clear the preserved cache before reboot using any of the virtual disks present on the controller.

Security key errors

Secured foreign import errors

A foreign configuration is a RAID configuration that already exists on a replacement physical disk that you install in a system. A secured foreign configuration is a RAID configuration that was created under a different security key.

There are two scenarios in which a secured foreign import fails:

- The passphrase authentication fails—A virtual disk secured with a security key different from the current controller security key cannot be imported without authentication of the original passphrase used to secure them. Supply the correct passphrase to import the secured foreign configuration. If you have lost or forgotten the passphrase, the secured foreign disks remain locked (inaccessible) until the appropriate passphrase is entered or if they are erased.
- The secured virtual disk is in an offline state after supplying the correct passphrase—You must check to determine why the virtual disk failed and correct the problem.

Failure to select or configure non Self-Encrypting Disks non-SED

A virtual disk can be either secured or unsecured depending on how it was configured when created. In order to create a secured virtual disk, the controller must have a security key present and must contain SEDs only. In order to select/configure non-SED, you must create an unsecured virtual disk. You can create an unsecured virtual disk even if there is a security key. Select the **Secure VD** option as **No** in the **Create New VD** menu. For steps on how to create an unsecured virtual disk, see Create virtual disks.



Failure to delete security key

A security key is used to lock or unlock access to a security-enabled component. This key is not utilized in the actual encryption of data. If a security key is present, both secured and unsecured virtual disks may exist.

To delete the security key, you must have a previously established security key present on the controller and there cannot be any configured secured disks. If there are configured secured virtual disks, remove or delete them.

Failure of Cryptographic Erase on encryption-capable physical disks

Cryptographic Erase is the process of securely erasing all data permanently on an encryption-capable physical disk and resetting the security attributes. It is used in scenarios such as deleting a foreign configuration in the event of a forgotten or lost passphrase or unlocking a disk that had been previously locked.

You can perform Cryptographic Erase only on encryption-capable disks that are not hot spares and not configured as non-RAID or virtual disks. Ensure that the conditions are met and see Cryptographic Erase.

General issues

PERC card has yellow bang in Windows operating system device manager

Issue:	The device is displayed in Device Manager but has a yellow bang (exclamation mark).
Corrective Action:	Reinstall the driver. For more information on reinstalling drivers, see Driver support for PERC 11.

PERC card not seen in operating systems

Issue:	The device does not appear in the Device Manager
Corrective	Power off the system and reseat the controller.

Issues in controller, battery, and disk when operating at low temperature

lssue:	If the controller is operating at temperatures less than zero degree Centigrade, then an increase in the number of issues related to controller, battery, or drive is observed.
Corrective Action:	Ensure that the controller ambient temperature is more than zero degree Centigrade.

Physical disk issues

Physical disk in failed state

Issue:

One of the physical disks in the disk array is in the failed state.

92 Troubleshooting issues in PERC11 cards



Corrective Update the PERC cards to the latest firmware available on the support site and replace the drive. **Action:**

Unable to rebuild a fault tolerant virtual disk

Issue:	Cannot rebuild a fault tolerant virtual disk. For more information, see the alert log for virtual disks.
Probable Cause:	The replacement disk is too small or not compatible with the virtual disk.
Corrective Action:	Replace the failed disk with a compatible good physical disk with equal or greater capacity.

Fatal error or data corruption reported

Issue:	$\label{eq:Fatalerror} \ensuremath{Fatal}\xspace \ensuremath{error}(s) \mbox{ or data corruption}(s) \mbox{ are reported when accessing virtual disks}.$
Corrective Action:	Contact Global Technical Support.

Multiple disks are inaccessible

Issue:	Multiple disks are simultaneously inaccessible.
Probable Cause:	Multiple physical disk errors in a single array typically indicate a failure in cabling or connection and could involve the loss of data.
Corrective Action:	You can recover the virtual disk after multiple physical disks become simultaneously inaccessible. Perform the following steps to recover the virtual disk:

CAUTION: Follow the safety precautions to prevent electrostatic discharge.

- 1. Turn off the system, check cable connections, and reseat physical disks.
- 2. Ensure that all the disks are present in the enclosure.
- 3. Turn on the system and enter the HII Configuration Utility.
- 4. Import the foreign configuration.
- 5. Press <F> at the prompt to import the configuration, or press <C> to enter the **HII Configuration Utility** and either import or clear the foreign configuration.

If the virtual disk is redundant and transitioned to **Degraded** state before going **Offline**, a rebuild operation starts automatically after the configuration is imported. If the virtual disk has gone directly to the **Offline** state due to a cable pull or power loss situation, the virtual disk is imported in its **Optimal** state without a rebuild occurring.

NOTE: You can use the **HII Configuration Utility** or Dell OpenManage storage management application to perform a manual rebuild of multiple physical disks.

Rebuilding data for a failed physical disk

Issue:	Rebuilding data for a physical disk that is in a failed state.
Probable Cause:	Physical disk is failed or removed.
Corrective Action:	If you have configured hot-spares, the PERC card automatically tries to use one of the hot-spares to rebuild a physical disk that is in a failed state. Manual rebuild is necessary if no hot-spares with enough capacity to rebuild the failed physical disks are available. You must insert a physical disk with enough storage in the subsystem before rebuilding the physical disk.

NOTE: You can use the **HII Configuration Utility** or Dell OpenManage storage management application to perform a manual rebuild of an individual physical disk.



Virtual disk fails during rebuild using a global hot spare

Issue:	A virtual disk fails during rebuild while using a global hot spare.
Probable Cause:	One or more disks in the virtual disks fails or is disconnected while the rebuild is in progress.
Corrective Action:	No action is required. The global hot spare reverts to Hot spare state and the virtual disk is in Failed state.

Dedicated hot spare disk fails during rebuild

Issue:	A hot spare disk fails during rebuild while using a dedicated hot spare.
Probable Cause:	The dedicated hot spare assigned to the virtual disk fails or is disconnected while the rebuild is in progress.
Corrective Action:	If there is a global hot spare available with enough capacity, rebuild will automatically start on the global hot spare. Where there is no hot spare present, you must insert a physical disk with enough capacity into the system before performing a rebuild.

Redundant virtual disk fails during reconstruction

Issue:	Multiple disks fails during a reconstruction process on a redundant virtual disk that has a hot spare.
Probable Cause:	Multiple physical disks in the virtual disk is failed or the cables are disconnected.
Corrective Action:	No action is required. The physical disk to which a reconstruction operation is targeted reverts to Ready state, and the virtual disk goes to Failed state. If there are any other virtual disks that can be supported by the capacity of the hot spare then the dedicated hot spare is converted to global hot spare, if not the hot spare will revert back to Ready state.

Virtual disk fails rebuild using a dedicated hot spare

Issue:	A virtual disk fails during rebuild while using a dedicated hot spare.
Probable Cause:	One or more disks in the virtual disks fails or is disconnected while the rebuild is in progress.
Corrective Action:	No action is required. The dedicated hot spare is in hot spare state and converted to global hot spare if there is any other virtual disk that is supported, otherwise the dedicated hot spare reverts to Ready state and the virtual drive is in Failed state.

Physical disk takes a long time to rebuild

Issue:	A physical disk is taking longer than expected to rebuild.
Description:	A physical disk takes longer to rebuild when under high I/O stress. There is only one rebuild I/O operation for every five host I/O operations.
Corrective Action:	If possible, reduce I/O stress on the physical disk or increase the value of rebuild rate controller parameter.

Drive removal and insertion in the same slot generates a foreign configuration event

Issue:

When a drive which is part of a virtual disk is removed and reinserted into the same slot the drive goes through a transient state of being foreign for a short period of time before rebuilding.

94 Troubleshooting issues in PERC11 cards

Inserido ao protocolo **22.951.206-4** por: **Pedro Henrique Golin Linhares** em: 31/03/2025 12:49. A autenticidade deste documento pode ser validada no endereço: https://www.eprotocolo.pr.gov.br/spiweb/validarDocumento com o código: 692a5ef82beb173ed33dad4ee6505546.



Description: This transient state could be reported as an event in management applications as A foreign configuration was detected on RAID Controller is SL x, where x is the slot of the RAID controller.

CorrectiveNo action is required on the foreign configuration state of the drive as it is transient and the controllerAction:handles the event automatically.

SMART errors

SMART monitors the internal performance of all motors, heads, and physical disk electronics and detects predictable physical disk failures.

NOTE: For information about SMART errors' reports that could indicate hardware failure, see the *Dell OpenManage Storage Management User's Guide* available at OpenManage Manuals.

Smart error detected on a non-RAID disk

Issue:

A SMART error is detected on a non-RAID disk.

Corrective Action:

- Perform the following steps: **1.** Back up your data.
 - Back up your data.
 Replace the affected physical disk with a new physical disk of equal or higher capacity.
 - **3.** Restore from the backup.

Smart error detected on a physical disk in a non-redundant virtual disk

lssue:	A SMART error is detected on a physical disk in a non-redundant virtual disk.
Corrective	Perform the following steps:
Action:	1. Back up your data.
	2. Use Replace Member to replace the disk manually.

(i) **NOTE:** For more information about the **Replace Member** feature, see Configure hot spare drives.

- 3. Replace the affected physical disk with a new physical disk of equal or higher capacity.
- **4.** Restore from the backup.

Smart error detected on a physical disk in a redundant virtual disk

lssue:	A SMART error is detected on a physical disk in a redundant virtual disk.
Corrective Action:	Perform the following steps:1. Back up your data.2. Force the physical disk offline.

(i) **NOTE:** If a hot spare is present, the rebuild starts with the hot spare after the disk is forced offline.

- 3. Replace the disk with a new physical disk of equal or higher capacity.
- 4. Perform the **Replace Member** operation.

() NOTE: The **Replace Member** operation allows you to copy data from a source physical disk of a virtual disk to a target physical disk that is not a part of the virtual disk. For more information about the **Replace Member** feature, see the topic Configure hot spare drives.

Troubleshooting issues in PERC11 cards 95



Replace member errors

(i) NOTE: For more information about the **Replace Member** features, see Configure hot spare drives.

Source disk fails during replace member operation

lssue:	The source disk fails during the Replace Member operation and the Replace Member operation stops due to the source physical disk error.
Probable Cause:	Physical disk failure or physical disk is removed or disconnected.
Corrective Action:	No action required. If the virtual disk can tolerate disk failure, and the source data is available from other disks in the virtual disk, the rebuild begins automatically on the target disk, using the data from the other disks, if the virtual disk cannot tolerate the failure, the virtual disk goes to offline state and the replace member operation is stopped.

Target disk fails during replace member operation

Issue:	The target disk failure reported during the Replace Member operation, and the Replace Member operation stops.
Probable Cause:	Physical disk failure or physical disk is removed or disconnected.
Corrective Action:	It is recommended that you replace or check the target drive, and restart the Replace Member operation or perform the operation on a different target drive.

A member disk failure is reported in the virtual disk which undergoes replace member operation

Issue:The source and the target drive which is part of Replace Member operation are online, while a different
drive which is a member of the virtual drive reports a failure.Probable Cause:Physical disk failure or physical disk is removed or disconnected.Corrective
Action:A rebuild starts if there any hot-spares configured or you may replace the failed drive. The Replace
Member operation continues as far as the source virtual disk can tolerate the drive failure. If the source
virtual disk fails, the Replace Member is stopped, otherwise the virtual disk continues to be in degraded
state

Linux operating system errors

Virtual disk policy is assumed as write-through

Error:<Date:Time> <HostName> kernel: sdb: asking for cache data failed<Date:Time> <HostName> kernel: sdb: assuming drive cache: write throughCorrective
Action:The error message is displayed when the Linux Small Computer System Interface (SCSI) mid-layer asks for physical disk cache settings. The controller firmware manages the virtual disk cache settings on a per controller and a per virtual disk basis, so the firmware does not respond to this command. The Linux SCSI mid-layer assumes that the virtual disk's cache policy is Write-Through. SDB is the device node for a virtual disk. This value changes for each virtual disk. Except for this message, there is no effect of this behavior on normal operation. The cache policy of the virtual disk and the I/O throughput are not affected by this message. The cache policy settings for the PERC SAS RAID system remain unchanged.

Inserido ao protocolo **22.951.206-4** por: **Pedro Henrique Golin Linhares** em: 31/03/2025 12:49. A autenticidade deste documento pode ser validada no endereço: https://www.eprotocolo.pr.gov.br/spiweb/validarDocumento com o código: 692a5ef82beb173ed33dad4ee6505546.


Unable to register SCSI device error message

Error:

smartd[smartd[2338] Device: /dev/sda, Bad IEC (SMART) mode page, err=-5, skip device smartd[2338] Unable to register SCSI device /dev/sda at line 1 of file /etc/smartd.conf.

CorrectiveThis is a known issue. An unsupported command is entered through the user application. User applications
attempt to direct Command Descriptor Blocks to RAID volumes. The error message does not affect the
feature functionality. The Mode Sense/Select command is supported by firmware on the controller.
However, the Linux kernel daemon issues the command to the virtual disk instead of to the driver IOCTL
node. This action is not supported.

Drive indicator codes

The LEDs on the drive carrier indicates the state of each drive. Each drive carrier has two LEDs: an activity LED (green) and a status LED (bicolor, green/amber). The activity LED blinks whenever the drive is accessed.



Figure 30. Drive indicators

- 1. Drive activity LED indicator
- 2. Drive status LED indicator
- 3. Drive capacity label

If the drive is in the Advanced Host Controller Interface (AHCI) mode, the status LED indicator does not power on. Drive status indicator behavior is managed by Storage Spaces Direct. Not all drive status indicators may be used.

Table 21. Drive indicator codes

Drive status indicator code	Condition
Blinks green twice per second	The drive is being identified or preparing for removal
Off	The drive is ready for removal i NOTE: The drive status indicator remains off until all drives are initialized after the system is powered on. Drives are not ready for removal during this time.
Blinks green, amber, and then powers off	There is an expected drive failure
Blinks amber four times per second	The drive has failed
Blinks green slowly	The drive is rebuilding
Solid green	The drive is online
Blinks green for three seconds, amber for three seconds, and then powers off after six seconds	The rebuild has stopped



HII error messages

Unhealthy Status of the drivers

Error:	One or more boot driver(s) have reported issues. Check the Driver Health Menu in Boot Manager for details.
Probable Cause:	This message might indicate that the cables are not connected, the disks might be missing, or the UEFI driver might require configuration changes.
Corrective Action:	 Check if the cables are connected properly, or replace missing hard drives, if any and then restart the system. Press any key to load the driver health manager to display the configurations. The Driver Health Manager displays the driver(s), which requires configuration. Alternately, if the UEFI driver requires configuration, press any key to load the Configuration Utility.

Rebuilding a drive during full initialization

Issue:	Automatic rebuild of drives is disabled for virtual disk during full initialization.
Corrective Action:	After full initialization the drive will automatically start its rebuild on its corresponding virtual disk.

System reports more drive slots than what is available

The system reports more slots than what is available in the following two scenarios:

System drives are hot swappable with backplane.	When the system drives are hot swappable, the PERC controller is not able to communicate correctly with the backplane or enclosure. Hence, the PERC controller reports a generic enclosure with drive 16 slots. In iDRAC, under Overview > Enclosures , the Enclosure ID is displayed as BP_PSV and Firmware version is displayed as 03 .
Corrective action	Turn off the system, reseat the controller and all the cables on the controller and backplane. If the issue is not resolved, contact your Dell Technical Service representative.
System drives are not hot swappable with cable direct attached.	When the system drives are not hot swappable, a default enclosure with 16 drive slots is expected to be reported (even though the system does not support that many drives).

World Wide Number on drive sticker is not the same in applications

World Wide Number (WWN) on the drive sticker and applications are not matching. NVMe drives do not have a WWN. So, the applications create a WWN from the available drive information. This WWN may not match with the WWN that is on the drive sticker, if present.

98 Troubleshooting issues in PERC11 cards



Backplane firmware revision not changing in PERC interfaces after an update

After updating the backplane firmware on 15G and later PowerEdge servers, the backplane version will not show as updated on some interfaces until the system is reset.

Troubleshooting issues in PERC11 cards 99



Appendix RAID description

RAID is a group of independent physical disks that provides high performance by increasing the number of disks used for saving and accessing data.

CAUTION: In the event of a physical disk failure, a RAID 0 virtual disk fails, resulting in data loss.

A RAID disk subsystem offers the following benefits:

- Improved I/O performance and data availability.
- Improved data throughput because several disks are accessed simultaneously. The physical disk group appears either as a single storage unit or multiple logical units to the host system.
- Improved data storage availability and fault tolerance. Data loss caused by a physical disk failure can be recovered by rebuilding missing data from the remaining physical disks containing data or parity.

Topics:

- Summary of RAID levels
- RAID 10 configuration
- RAID terminology

Summary of RAID levels

Following is a list of the RAID levels supported by the PERC 11 series of cards:

- RAID 0 uses disk striping to provide high data throughput, especially for large files in an environment that requires no data redundancy.
- RAID 1 uses disk mirroring so that data written to one physical disk is simultaneously written to another physical disk. RAID 1 is good for small databases or other applications that require small capacity and complete data redundancy.
- RAID 5 uses disk striping and parity data across all physical disks (distributed parity) to provide high data throughput and data redundancy, especially for small random access.
- RAID 6 is an extension of RAID 5 and uses an additional parity block. RAID 6 uses block-level striping with two parity blocks distributed across all member disks. RAID 6 provides protection against double disk failures, and failures while a single disk is rebuilding. If you are using only one array, deploying RAID 6 is more effective than deploying a hot spare disk.
- RAID 10 is a combination of RAID 0 and RAID 1, uses disk striping across mirrored disks. It provides high data throughput and complete data redundancy.
- RAID 50 is a combination of RAID 0 and RAID 5 where a RAID 0 array is striped across RAID 5 elements. RAID 50 requires at least six disks.
- RAID 60 is a combination of RAID 0 and RAID 6 where a RAID 0 array is striped across RAID 6 elements. RAID 60 requires at least eight disks.

The following table lists the minimum and maximum disks supported on each RAID levels.

Table 22. Minimum and maximum disks supported on each RAID levels

RAID Level	Minimum disk	Maximum disk
0	1	32
1	2	2
5	3	32
6	4	32
10	4	240
50	6	240
60	8	240



(i) NOTE: The maximum number of virtual disks is currently limited to 192, because of the supported enclosure configuration.

RAID 10 configuration

In PERC 10 and PERC 11 controllers, RAID 10 can be configured without spanning up to 32 drives. Any RAID 10 volume that has more than 32 drives require spanning. Each span can contain up to 32 drives. Drives must be distributed evenly across all the spans with each span containing an even number of drives.

NOTE: Spans in a RAID 10 volume are only supported if spans are even. Uneven spanned RAID 10 cannot be imported from previous controller generations.

The following table shows the RAID 10 configurations.

Table 23. RAID 10 configurations

Disk or span count	RAID 10 capable						
4 (1)	Yes	64 (2)	Yes	124	No	184	No
6 (1)	Yes	66 (3)	Yes	126 (7)	Yes	186	No
8 (1)	Yes	68	No	128 (4)	Yes	188	No
10 (1)	Yes	70 (5)	Yes	130 (5)	Yes	190	No
12 (1)	Yes	72 (3)	Yes	132 (6)	Yes	192 (6)	Yes
14 (1)	Yes	74	No	134	No	194	No
16 (1)	Yes	76	No	136	No	196 (7)	Yes
18 (1)	Yes	78 (3)	Yes	138	No	198	No
20 (1)	Yes	80 (4)	Yes	140 (5)	Yes	200	No
22 (1)	Yes	82	No	142	No	202	No
24 (1)	Yes	84 (6)	Yes	144	Yes	204	No
26 (1)	Yes	86	No	146	No	206	No
28 (1)	Yes	88 (4)	Yes	148	No	208 (8)	Yes
30 (1)	Yes	90 (3)	Yes	150 (5)	Yes	210 (7)	Yes
32 (1)	Yes	92	No	152	No	212	No
34	No	94	No	154 (7)	Yes	214	No
36 (2)	Yes	96 (3)	Yes	156 (6)	Yes	216	No
38	No	98 (7)	Yes	158	No	218	No
40 (2)	Yes	100 (5)	Yes	160 (5)	Yes	220	No
42 (2)	Yes	102	No	162	No	222	No
44 (2)	Yes	104 (4)	Yes	164	No	224 (8)	Yes
46	No	106	No	166	No	226	No
48 (2)	Yes	108 (6)	Yes	168 (6)	Yes	228	No
50 (2)	Yes	110 (5)	Yes	170	No	230	No
52 (2)	Yes	112 (4)	Yes	172	No	232	No
54 (2)	Yes	114	No	174	No	234	No
56 (2)	Yes	116	No	176 (8)	Yes	236	No
58	No	118	No	178	No	238	No



Table 23. RAID 10 configurations (continued)

Disk or span count	RAID 10 capable						
60 (2)	Yes	120 (4)	Yes	180 (6)	Yes	240 (8)	Yes
62	No	122	No	182 (7)	Yes	-	-

RAID terminology

Disk striping

Disk striping allows you to write data across multiple physical disks instead of just one physical disk. Disk striping involves partitioning each physical disk storage space in stripes of the following sizes: 64 KB, 128 KB, 256 KB, 512 KB, and 1 MB. The stripes are interleaved in a repeated sequential manner. The part of the stripe on a single physical disk is called a stripe element.

For example, in a four-disk system using only disk striping (used in RAID 0), segment 1 is written to disk 1, segment 2 is written to disk 2, and so on. Disk striping enhances performance because multiple physical disks are accessed simultaneously, but disk striping does not provide data redundancy.



Figure 31. Example of disk striping (RAID 0)

Disk mirroring

With mirroring (used in RAID 1), data written to one disk is simultaneously written to another disk. If one disk fails, the contents of the other disk can be used to run the system and rebuild the failed physical disk. The primary advantage of disk mirroring is that it provides complete data redundancy. Both disks contain the same data at all times. Either of the physical disks can act as the operational physical disk.

Disk mirroring provides complete redundancy, but is an expensive option because each physical disk in the system must be duplicated.





Stripe element 2 Stripe element 2 Duplicated Stripe element 3 Stripe element 3 Duplicated Stripe element 3 Stripe element 4 Duplicated

Figure 32. Example of Disk Mirroring (RAID 1)



Spanned RAID levels

Spanning is a term used to describe the way in which RAID levels 10, 50, and 60 are constructed from multiple sets of basic, or simple RAID levels. For example, a RAID 10 has multiple sets of RAID 1 arrays where each RAID 1 set is considered a span. Data is then striped (RAID 0) across the RAID 1 spans to create a RAID 10 virtual disk. Similarly, RAID 50 and RAID 60 combine multiple sets of RAID 5 or RAID 6 respectively with striping.

Parity data

Parity data is redundant data that is generated to provide fault tolerance within certain RAID levels. In the event of a disk failure, the parity data can be used by the controller to regenerate user data. Parity data is present for RAID 5, 6, 50, and 60.

The parity data is distributed across all the physical disks in the system. If a single physical disk fails, it can be rebuilt from the parity and the data on the remaining physical disks. RAID level 5 combines distributed parity with disk striping. Parity provides redundancy for one physical disk failure without duplicating the contents of the entire physical disks.

RAID 6 combines dual distributed parity with disk striping. This level of parity allows for two disk failures without duplicating the contents of entire physical disks.



Strpe element 1	Stripe element 2	Stripe element 3.	Stipe element 4	Shipe element 5	Parity (t-6)
Shipo element 7	Stripe element 8	Stripe element 9	Stripe element 10	Parity (6-10)	Strips element 6
Stripe element 13	Stripe element 14	Strips element 15	Parity (11-15)	Stripe element 11	Shipe element 12
Stripe element 19	Stripe element 20	Parity (16-20)	Stripe element 16	Shipe element 17	Stripe element 18
Stripe element 25	Parity (21-25)	Stripe element 21	Stripe element 22	Siripe element 23	Stripe element 24
Parity (26-30)	Stripe element 28	Stripe element 27	Stripe element 28	Stripe-element 29	Stripe element 30

Figure 33. Example of Distributed Parity (RAID 5)

(i) NOTE: Parity is distributed across multiple physical disks in the disk group.



Figure 34. Example of Dual Distributed Parity (RAID 6)

(i) NOTE: Parity is distributed across all disks in the array.



Getting help

Topics:

- Recycling or End-of-Life service information
- Contacting Dell
- Locating the Express Service Code and Service Tag
- Receiving automated support with SupportAssist

Recycling or End-of-Life service information

Take back and recycling services are offered for this product in certain countries. If you want to dispose of system components, visit the How to Recycle page and select the relevant country.

Contacting Dell

Dell provides online and telephone based support and service options. If you do not have an active internet connection, you can find Dell contact information on your purchase invoice, packing slip, bill or Dell product catalog. The availability of services varies depending on the country and product, and some services may not be available in your area. To contact Dell for sales, technical assistance, or customer service issues:

Steps

- 1. Go to the Support site.
- 2. Select your country from the drop-down menu on the lower right corner of the page.
- **3.** For customized support:
 - a. Enter the system Service Tag in the Enter a Service Tag, Serial Number, Service Request, Model, or Keyword field.
 - b. Click Submit.

The support page that lists the various support categories is displayed.

- 4. For general support:
 - **a.** Select your product category.
 - b. Select your product segment.
 - c. Select your product.
 - The support page that lists the various support categories is displayed.
- **5.** For contact details of Dell Global Technical Support:
 - a. Click Global Technical Support.
 - b. The **Contact Technical Support** page is displayed with details to call, chat, or e-mail the Dell Global Technical Support team.

Locating the Express Service Code and Service Tag

The unique Express Service Code and Service Tag is used to identify the system.

The information tag is located on the front of the system rear of the system that includes system information such as Service Tag, Express Service Code, Manufacture date, NIC, MAC address, QRL label, and so on. If you have opted for the secure default access to iDRAC, the Information tag also contains the iDRAC secure default password. If you have opted for iDRAC Quick Sync 2, the Information tag also contains the OpenManage Mobile (OMM) label, where administrators can configure, monitor, and troubleshoot the PowerEdge servers.





Figure 35. Locating the Express Service Code and Service tag

- 1. Information tag (front view)
- 3. OpenManage Mobile (OMM) label
- 5. Service Tag, Express Service Code, QRL label
- 2. Information tag (back view)
- 4. iDRAC MAC address and iDRAC secure password label

The Mini Enterprise Service Tag (MEST) label is located on the rear of the system that includes Service Tag (ST), Express Service Code (Exp Svc Code), and Manufacture Date (Mfg. Date). The Exp Svc Code is used by Dell to route support calls to the appropriate personnel.

Alternatively, the Service Tag information is located on a label on left wall of the chassis.

Receiving automated support with SupportAssist

Dell SupportAssist is an optional Dell Services offering that automates technical support for your Dell server, storage, and networking devices. By installing and setting up a SupportAssist application in your IT environment, you can receive the following benefits:

- Automated issue detection SupportAssist monitors your Dell devices and automatically detects hardware issues, both proactively and predictively.
- Automated case creation When an issue is detected, SupportAssist automatically opens a support case with Dell Technical Support.
- Automated diagnostic collection SupportAssist automatically collects system state information from your devices and uploads it securely to Dell. This information is used by Dell Technical Support to troubleshoot the issue.
- Proactive contact A Dell Technical Support agent contacts you about the support case and helps you resolve the issue.

The available benefits vary depending on the Dell Service entitlement purchased for your device. For more information about SupportAssist, go to the SupportAssist page.



Documentation resources

This section provides information about the documentation resources for your system.

To view the document that is listed in the documentation resources table:

- From the Dell support site:
 - 1. Click the documentation link that is provided in the Location column in the table.
 - 2. Click the required product or product version.

(i) NOTE: To locate the product name and model, see the front of your system.

3. On the Product Support page, click Manuals & documents.

- Using search engines:
 - Type the name and version of the document in the search box.

Table 24. Additional documentation resources for your system

Task	Document	Location
Setting up your system	For more information about installing and securing the system into a rack, see the Rail Installation Guide included with your rail solution. For information about setting up your system, see the Getting Started Guide document that is shipped with your system	PowerEdge Server Manuals
Configuring your system	For information about the iDRAC features, configuring and logging in to iDRAC, and managing your system remotely, see the Integrated Dell Remote Access Controller User's Guide.	PowerEdge Server Manuals
	For information about understanding Remote Access Controller Admin (RACADM) subcommands and supported RACADM interfaces, see the RACADM CLI Guide for iDRAC.	
	For information about Redfish and its protocol, supported schema, and Redfish Eventing implemented in iDRAC, see the Redfish API Guide.	
	For information about iDRAC property database group and object descriptions, see the Attribute Registry Guide.	
	For information about Intel QuickAssist Technology, see the Integrated Dell Remote Access Controller User's Guide.	
	For information about earlier versions of the iDRAC documents.	iDRAC Manuals
	To identify the version of iDRAC available on your system, on the iDRAC web interface, click	



Table 24. Additional documentation resources for your system (continued)

Task	Document	Location
	? > About.	
	For information about updating drivers and firmware, see the Methods to download firmware and drivers section in this document.	Drivers
Understanding event and error messages	For information about the event and error messages generated by the system firmware and agents that monitor system components, go to qrl.dell.com > Look Up > Error Code, type the error code, and then click Look it up.	PowerEdge Server Event and Error Messages

Documentation resources 107



Home > Hardware > PowerEdge R760XS

By Dell Inc

PowerEdge R760XS

No description provided.

Product type	Hardware
Processor architecture	X64
Vendor	Dell Inc

is a certification that ensures hardware, like servers and storage devices, meets specific reliability and quality standards. Read more about <u>Additional Qualification</u> <u>Certification</u>.

Windows	Product type	Hardware
erver 025	Manufacturer	Dell Inc.
Contifical	Processor name	INTEL(R) XEON(R) GOLD 6534
Lertined	Tested memory	1536GB 5600MHz

ls. 858

 \mathcal{O}

- > View Additional Qualifications and features
- > View and download submission(s)

Windows	Product type	Hardware
Server 2022	Manufacturer	Dell Inc.
	Processor name	INTEL(R) XEON(R) GOLD 6534
Certified	Tested memory	1536GB 5600MHz



D&LLTechnologies

Specification Sheet





PowerEdge R760xs

Best choice in balanced compute and flexible storage for the most popular IT applications

Buy the performance and flexibility you need

The new Dell PowerEdge R760xs is a 2U, two-socket rack server. Buy the best fit in scalable performance and large storage capability with this purpose-built 2U system. Focused on delivering the latest technology to power the most popular applications and workloads used by businesses today, including virtual desktop infrastructure (VDI), virtual machines (VMs), and software-defined storage (SDS). All delivered in a thoughtfully crafted platform that will provide balanced compute that fits in your current infrastructure.

Easily configurable

- Add up to two 4th generation Intel® Xeon® Scalable processors with up to 32 cores per socket for faster performance
- · Accelerate in-memory workloads with up to 16 DDR5 RDIMMS up to 4800 MT/sec
- Improve data throughput and reduce latency with support up to 8 I/O device (up to available 6 PCIe slots, 1 OCP 3.0 networking slot, and 1 dedicated PERC slot)
- Storage options include up to 12x 3.5" HDDs/SSDs, or up to 16x 2.5" HDD/SSDs, plus up to 8x NVMe drives

A breeze to cool

- Thoughtfully designed to fit in your current air-cooled infrastructure
- · Alleviate the worry about expensive liquid cooling retrofitting to your data center
- Synchronize your workload needs with a tailored performance configuration that is air cooled
- Minimize the carbon footprint of your data center by better matching the system power consumption with anticipated workload requirements

Cyber Resilient Architecture for Zero Trust IT environment & operations

Security is integrated into every phase of the PowerEdge lifecycle, including protected supply chain and factory-to-site integrity assurance. Silicon-based root of trust anchors end-to-end boot resilience while Multi-Factor Authentication (MFA) and role-based access controls ensure trusted operations.

Increase efficiency and accelerate operations with autonomous collaboration

The Dell OpenManage[™] systems management portfolio delivers a secure, efficient, and comprehensive solution for PowerEdge servers. Simplify, automate and centralize one-to-many management with the OpenManage Enterprise console and iDRAC.

Sustainability

From recycled materials in our products and packaging, to thoughtful, innovative options for energy efficiency, the PowerEdge portfolio is designed to make, deliver, and recycle products to help reduce the carbon footprint and lower your operation costs. We even make it easy to retire legacy systems responsibly with Dell Technologies Services.

Rest easier with Dell Technologies Services

Maximize your PowerEdge Servers with comprehensive services ranging from Consulting, to ProDeploy and ProSupport suites, Data Migration and more – available across 170 locations and backed by our 60K+employees and partners.

PowerEdge R760xs

The Dell PowerEdge R760xs offers compelling performance in a right-sized system with the latest PCIe Gen 5 bandwidth and large storage capability to support:

- Virtual Desktop Infrastructure (VDI)
- Virtual Machines (VMs)
- Software-Defined Storage Node

Fasture				
Feature	Technical Specifications			
Processor	Up to two 4th Generation Intel Xeon Scalable processor with up to 32 cores per processor			
Memory	 16 DDR5 DIMM slots, supports RDIMM 1 TB max, speeds up to 4800 MT/s, supports registered ECC DDR5 DIMMs only 			
Storage controllers	 Internal Controllers: PERC H965i, PERC H755, PERC H755N, PE Internal Boot: Boot Optimized Storage Subsystem (BOSS-N1): HV External HBA (non-RAID): HBA355e; Software RAID: S160 	:RC H355, HBA355i VRAID 1, 2 x M.2 NVMe SSDs or USB		
GPU Options	2 x 75 W SW, LP			
Drive Bays	 Front bays: 0 drive bay Up to 8 x 3.5-inch SAS/SATA (HDD/SSD) max 160 TB Up to 12 x 3.5-inch SAS/SATA (HDD/SSD) max 240 TB Up to 8 x 2.5-inch SAS/SATA/NVMe (HDD/SSD) max 122.88 TB Up to 16 x 2.5-inch SAS/SATA (HDD/SSD) max 121.6 TB Up to 16 x 2.5-inch (SAS/SATA) + 8 x 2.5-inch (NVMe) (HDD/SSD) max 244.48 TB 	 Rear bays: Up to 2 x 2.5-inch SAS/SATA/NVMe (HDD/SSD) max 30.72 TB (supported only with 12 x 3.5-inch SAS/SATA HDD/SSD configuration) 		
Hot swap Redundant Power Supplies	 1800 W Titanium 200—240 VAC or 240 HVDC 1400 W Platinum 100—240 VAC or 240 HVDC 1100 W Titanium 100—240 VAC or 240 HVDC 1100 W LVDC -48 — (-60) VDC 	 800 W Platinum 100—240 VAC or 240 HVDC, 700 W Titanium 200—240 VAC or 240 HVDC 600 W Platinum 100—240 VAC or 240 HVDC 		
Cooling Options	Air cooling			
Fans	• Standard (STD) fans/High performance Silver (HPR) fans/ High per	erformance Gold (VHP) fans, Up to 6 hot swappable fans		
Dimensions and Weight	 Height – 86.8 mm (3.41 inches) Width – 482 mm (18.97 inches) 	 Depth – 707.78 mm (27.85 inches) – without bezel 721.62 mm (28.4 inches) – with bezel Weight – Max 28.6 kg (63.0 lbs.) 		
Form Factor	2U rack server			
Embedded Management	 iDRAC9 iDRAC Direct iDRAC RESTful API with Redfish 	iDRAC Service ModuleQuick Sync 2 wireless module		
Bezel	Optional LCD bezel or security bezel			
OpenManage Software	 CloudIQ for PowerEdge plug in OpenManage Enterprise OpenManage Enterprise Integration for VMware vCenter OpenManage Integration for Microsoft System Center 	 OpenManage Integration with Windows Admin Center OpenManage Power Manager plugin OpenManage Service plugin OpenManage Update Manager plugin 		
Mobility	OpenManage Mobile			
OpenManage Integrations	 BMC Truesight Microsoft System Center OpenManage Integration with ServiceNow 	 Red Hat Ansible Modules Terraform Providers VMware vCenter and vRealize Operations Manager 		
Security	 Cryptographically signed firmware Data at Rest Encryption (SEDs with local or external key mgmt) Secure Boot Secure Erase 	 Secured Component Verification (Hardware integrity check) Silicon Root of Trust System Lockdown (requires iDRAC9 Enterprise or Datacenter) TPM 2.0 FIPS, CC-TCG certified, TPM 2.0 China NationZ 		
Embedded NIC	2 x 1 GbE LOM			
Network options	1 x OCP card 3.0 (optional)			
Ports	Front Ports: • 1 x iDRAC Direct (Micro-AB USB) port, 1 x USB 2.0, 1 x VGA Internal Ports: 1 x USB 3.0 (optional)	 Rear Ports 1 x Dedicated iDRAC Ethernet port, 1 x USB 2.0, 1 x USB 3.0, 1 x VGA, 1 x Serial (optional) 		
PCle	 1 CPU Configuration: Up to 4 PCIe slots (2 x8 Gen5, 1 x16 Gen4, 2 CPU configuration: Up to 6 PCIe slots (2 x16 Gen5, 3 x16 Gen4 	1 x8 Gen4) , 1 x8 Gen4)		
Operating System and Hypervisors	Microsoft Windows Server with Hyper-V Red Hat Enterprise Linux SUSE Linux Enterprise Server	VMware ESXi Canonical Ubuntu Server LTS For apopiliations and interpretability datails and Dall apr (00)		
OEM-ready version available	 SOSE LINUX Enterprise Server From bezel to BIOS to packaging, your servers can look and feel as if the Solutions -> OEM Solutions. 	they were designed and built by you. For more information, visit Dell.com ->		

APEX Flex on Demand

Acquire the technology you need to support your changing business with payments that scale to match actual usage. For more information, visit https://www.delltechnologies.com/en-us/payment-solutions/flexible-consumption/flex-on-demand.htm.



Copyright © 2023 Dell Inc. or its subsidiaries. All Rights Reserved. Dell Technologies, Dell, and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be trademarks of their respective owners.

D&LLTechnologies

April 2023



Dell PowerEdge R760xs

Technical Guide

Regulatory Model: E88S Regulatory Type: E88S001 November 2024 Rev. A03





Notes, cautions, and warnings

(i) NOTE: A NOTE indicates important information that helps you make better use of your product.

CAUTION: A CAUTION indicates either potential damage to hardware or loss of data and tells you how to avoid the problem.

WARNING: A WARNING indicates a potential for property damage, personal injury, or death.

© 2023 Dell Inc. or its subsidiaries. All rights reserved. Dell Technologies, Dell, and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be trademarks of their respective owners.



Chapter 1: System overview	5
Key workloads	5
New technologies	5
Chapter 2: System features and generational comparison	7
Chapter 3: Chassis views and features	9
Chassis views	9
Front view of the system	9
Rear view of the system	12
Inside the system	14
System diagnostics and indicator codes	15
QR code for PowerEdge R760xs system resources	18
Chapter 4: Processor	19
· Processor features	19
Supported processors	19
Chipset	20
Chipset features	20
Chapter 5: Memory subsystem	21
Supported memory	21
General memory module installation guidelines	21
Chapter 6: Storage	23
Storage controllers	23
Storage controller feature matrix	23
Server storage controllers User Guide	25
RAID - Redundant Array of Independent Disks	25
Datasheets and PERC performance scaling decks	25
Boot Optimized Storage Solution (BOSS)	25
Supported Drives	27
External Storage	27
Chapter 7: Networking	. 28
Overview	28
OCP 3.0 support	28
Supported OCP cards	28
OCP NIC 3.0 vs. rack Network Daughter Card comparisons	29
OCP form factors	29
Chapter 8: Slot priority matrix	31
Expansion card installation guidelines	31



Chapter 9: Power, thermal, and acoustics	
Power	
PSU specifications	
Thermal	41
Acoustics	41
Acoustical configurations of R760xs	41
Chapter 10: Rack, rails, and cable management	
Rails information	
A11 Sliding Rails features summary	43
A8 Static Rails features summary	
Cable Management Arm	
Strain Relief Bar	
Rack Installation	46
Chapter 11: Supported operating systems	51
Chapter 12: Doll OpenManage Systems Management	52
Integrated Dell Remote Access Controller (iDRAC)	52 52
Systems Management software support matrix	53 53
Chanter 13: Annendix A. Standards compliance	55
Chapter 14: Appendix B: Additional specifications	
Chapter 14: Appendix B: Additional specifications Chassis dimensions	56
Chapter 14: Appendix B: Additional specifications Chassis dimensions System weight	56
Chapter 14: Appendix B: Additional specifications Chassis dimensions System weight Video specifications	
Chapter 14: Appendix B: Additional specifications Chassis dimensions System weight Video specifications USB ports specifications	56
Chapter 14: Appendix B: Additional specifications Chassis dimensions System weight Video specifications USB ports specifications Environmental specifications	56 56 57 57 57 58
Chapter 14: Appendix B: Additional specifications Chassis dimensions System weight Video specifications USB ports specifications Environmental specifications Thermal restriction matrix	56 56 57 57 57 58 58 58
Chapter 14: Appendix B: Additional specifications Chassis dimensions System weight Video specifications USB ports specifications Environmental specifications Thermal restriction matrix Chapter 15: Appendix C Additional resources	56 56 57 57 57 58 58 60
Chapter 14: Appendix B: Additional specifications. Chassis dimensions System weight Video specifications. USB ports specifications. Environmental specifications. Thermal restriction matrix. Chapter 15: Appendix C Additional resources. Documentation.	56 56 57 57 58 58 60 64
Chapter 14: Appendix B: Additional specifications Chassis dimensions System weight Video specifications USB ports specifications Environmental specifications Thermal restriction matrix Chapter 15: Appendix C Additional resources Documentation Customer kits	56 57 57 58 58 58 60 64 65 65
Chapter 14: Appendix B: Additional specifications Chassis dimensions System weight Video specifications USB ports specifications Environmental specifications Thermal restriction matrix Chapter 15: Appendix C Additional resources Documentation Customer kits Dell Upgrades	56 56 57 57 58 58 60 64 65 65 65
Chapter 14: Appendix B: Additional specifications. Chassis dimensions System weight Video specifications. USB ports specifications. Environmental specifications. Thermal restriction matrix. Chapter 15: Appendix C Additional resources. Documentation. Customer kits Dell Upgrades Upgrades portfolio.	56 56 57 57 58 58 60 64 65 65 65 65 65
Chapter 14: Appendix B: Additional specifications Chassis dimensions System weight Video specifications USB ports specifications Environmental specifications Thermal restriction matrix Chapter 15: Appendix C Additional resources Documentation Customer kits Dell Upgrades Upgrades portfolio Upgrades reference links	56 56 57 57 58 58 58 60 64 65 65 65 65 65 65 65
Chapter 14: Appendix B: Additional specifications Chassis dimensions System weight Video specifications USB ports specifications Environmental specifications Thermal restriction matrix Chapter 15: Appendix C Additional resources Documentation Customer kits Dell Upgrades Upgrades portfolio Upgrades reference links Chapter 16: Appendix D: Service and support	56 56 57 57 58 58 60 64 65 65 65 65 65 65 65 65 65 65
Chapter 14: Appendix B: Additional specifications Chassis dimensions System weight Video specifications USB ports specifications Environmental specifications Thermal restriction matrix Chapter 15: Appendix C Additional resources Documentation Customer kits Dell Upgrades Upgrades portfolio Upgrades reference links Chapter 16: Appendix D: Service and support Why attach service contracts	56 56 57 57 58 58 60 64 65 65 65 65 65 65 67 68
Chapter 14: Appendix B: Additional specifications Chassis dimensions System weight Video specifications USB ports specifications Environmental specifications Thermal restriction matrix Chapter 15: Appendix C Additional resources Documentation Customer kits Dell Upgrades Upgrades portfolio Upgrades reference links Chapter 16: Appendix D: Service and support Why attach service contracts ProSupport Infrastructure Suite	56 56 57 57 58 58 58 58 58 58 58 58 58 58 58 58 58
Chapter 14: Appendix B: Additional specifications Chassis dimensions	56 56 57 57 58 58 58 60 64 65 65 65 65 65 65 65 65 65 65 65 70 70
Chapter 14: Appendix B: Additional specifications. Chassis dimensions System weight. Video specifications. USB ports specifications. Environmental specifications. Thermal restriction matrix. Chapter 15: Appendix C Additional resources. Documentation. Customer kits. Dell Upgrades. Upgrades portfolio. Upgrades reference links Chapter 16: Appendix D: Service and support. Why attach service contracts. ProSupport Infrastructure Suite Specialty Support Services. Dell deployment services.	56 56 57 57 58 58 58 60 64 65 65 65 65 65 65 65 65 67 67 70 70
Chapter 14: Appendix B: Additional specifications Chassis dimensions System weight	56 56 57 57 58 58 58 60 64 60 64 65 65 65 65 65 65 65 65 67 68 68 68 70 70
Chapter 14: Appendix B: Additional specifications Chassis dimensions System weight Video specifications USB ports specifications Environmental specifications Thermal restriction matrix Chapter 15: Appendix C Additional resources Documentation Customer kits Dell Upgrades Upgrades portfolio Upgrades reference links Chapter 16: Appendix D: Service and support Why attach service contracts ProSupport Infrastructure Suite Specialty Support Services Dell deployment services Dell deployment Scenarios	56 56 57 57 58 58 58 60 64 60 64 65 65 65 65 65 65 65 65 65 65 70 70 71
Chapter 14: Appendix B: Additional specifications Chassis dimensions System weight Video specifications USB ports specifications Environmental specifications Thermal restriction matrix Chapter 15: Appendix C Additional resources Documentation Customer kits Dell Upgrades Upgrades portfolio Upgrades reference links Chapter 16: Appendix D: Service and support Why attach service contracts ProSupport Infrastructure Suite Specialty Support Services Dell deployment services Dell deployment Services Dell deployment Services Dell deployment Services with Ansible DAY 2 – Automation Services with Ansible	56 57 57 58 58 60 64 65 65 65 65 65 65 65 65 65 65 70 70 71 75 76
Chapter 14: Appendix B: Additional specifications Chassis dimensions System weight Video specifications USB ports specifications Environmental specifications Thermal restriction matrix Chapter 15: Appendix C Additional resources Documentation Customer kits Dell Upgrades Upgrades portfolio Upgrades reference links Chapter 16: Appendix D: Service and support Why attach service contracts ProSupport Infrastructure Suite Specialty Support Services Dell deployment services Unique Deployment Services with Ansible DAY 2 – Automation Services with Ansible Dell Technologies Consulting Services	56 57 57 58 58 58 58 60 64 65 65 65 65 65 65 65 65 65 65 70 70 71 75 76 77



System overview

The Dell PowerEdge R760xs is Dell's latest two-socket, rack server that is designed to run complex workloads using highly scalable memory, I/O, and network options. The systems feature the 4thand 5th Gen Intel® Xeon® Scalable Processor (Socket E1/LGA4677-1), up to 16 DIMMs, PCI Express® (PCIe) 5.0 enabled expansion slots , and a broad selection of network interface technologies.

Topics:

- Key workloads
- New technologies

Key workloads

The target workloads for the PowerEdge R760xs include virtualization, medium density VM and VDI, and scale-out database. .

New technologies

Table 1. New technologies

Technology	Detailed Description
Up to 2 x 5th Gen Intel(R) Xeon(R) Scalable Processors (Emerald Rapids)	Core count: Up to 28 per processor
	Up to 3.9 GHz
	UPI speed: Up to 3x UPIs/Socket at 12.8 GT/s or 14.4 or 16 GT/s or 20 GT/s
	Maximum number of PCIe lanes: Integrated 80 PCIe 5.0 lanes @ 32 GT/s PCIe Gen5
	Maximum TDP: 250 W
Up to 2 x 4th Gen Intel(R) Xeon(R) Scalable Processors (Sapphire Rapids)	Core count: Up to 32 per processor
	UPI speed: Up to 3x UPIs/Socket at 12.8 GT/s or 14.4 or 16 GT/s
	Maximum number of PCIe lanes: Integrated 80 PCIe 5.0 lanes @ 32 GT/s PCIe Gen5
	Maximum TDP: 250 W
DDR5 ECC memory up to 5200 MT/s	Max 8 DIMMs per processor and 16 DIMMs per system
	Supports DDR5 ECC RDIMM
GPUs	Max 2 x 75 W SW GPUs (NVIDIA A2)
Flex I/O	LOM: 2x1GbE with BCM5720 LAN controller
	Rear I/O with: • 1x Dedicated iDRAC Ethernet port (1 GbE) • 1 x USB 3.0 • 1 x USB 2.0 • 1 x VGA port



Table 1. New technologies (continued)

Technology	Detailed Description
	Serial Port option
	Optional OCP Mezz 3.0 (supported by x8 PCIe lanes)
	Front I/O with: • 1 x USB 2.0 • 1x iDRAC Direct (Micro-AB USB) port • 1 x VGA port
CPLD 1-wire	Support payload data of Front PERC, Riser, BP and Rear IO to BOSS-N1 and iDRAC
Dedicated PERC	Front Storage module PERC with Front PERC11, PERC12 and PERC 12.2
Software RAID	OS RAID / S160
Power Supplies	60 mm dimension is the new PSU form factor on 16G design 600 W Platinum 100–240 VAC/ 240 VDC 700 W Titanium 200–240 VAC/ 240 VDC 800 W Platinum 100–240 VAC/ 240 VDC 1100 W DC/-48–(-60) V 1100 W Titanium 100–240 VAC/ 240 VDC 1400 W Titanium 100–240 VAC/ 240 VDC 1400 W Platinum 100–240 VAC/ 240 VDC 1400 W Titanium 277 VAC/ 336 VDC 1800 W Titanium 200–240 VAC/ 240 VDC



System features and generational comparison

The following table shows the comparison between the PowerEdge R760xs with the PowerEdge R750xs.

Table 2. Features comparison

Feature	PowerEdge R760xs	PowerEdge R750xs
Processor	 Up to 2 x 4th Gen Intel(R) Xeon(R) Scalable Processors (Sapphire Rapids) with up to 32 cores Up to 2 x 5th Gen Intel(R) Xeon(R) Scalable Processors (Emerald Rapids) with up to 28 cores 	Maximum two 3 rd Generation Intel [®] Xeon [®] Scalable processors with maximum 32 cores per processor
Processor Interconnect	Intel Ultra Path Interconnect (UPI) , up to 3 links per CPU	Intel Ultra Path Interconnect (UPI)
Memory	16 DDR5 DIMM slots	16 DDR4 DIMM slots
	Supports RDIMM 1.5 TB max	Supports RDIMM 1 TB max
	Speed maximum 5200 MT/s for 5th generation	Speed maximum 3200 MT/s
	and 4800 MT/s for 4th generation processors	Supports registered ECC DDR4 DIMMs only
	Supports registered ECC DDR5 DIMMs only	Apache Pass : No
	NVDIMM : No	NVDIMM : No
Storage Drives	 Front bays: 0 drive bay Maximum 8x 3.5-inch SAS/SATA (HDD/SSD) max 160 TB Maximum 12x 3.5-inch SAS/SATA (HDD/SSD) max 240 TB Maximum 8x 2.5-inch SAS/SATA/NVMe (HDD/SSD) max 122.88 TB Maximum 16x 2.5-inch SAS/SATA (HDD/SSD) max 121.6 TB Maximum 16x 2.5-inch (SAS/SATA) + 8x 2.5-inch (NVMe) (HDD/SSD) max 244.48 TB Rear bays: Maximum 2x 2.5-inch SAS/SATA (NVMe) 	 Front bays: 0 drive bay Maximum 8x 3.5-inch SAS/SATA (HDD/SSD) max 128 TB Maximum 12x 3.5-inch SAS/SATA (HDD/SSD) maximum 192 TB Maximum 8x 2.5-inch SAS/SATA/NVMe (HDD/ SSD) maximum 61.44 TB Maximum 16x 2.5-inch SAS/SATA (HDD/SSD) maximum 122.88 TB Maximum 16x 2.5-inch (SAS/SATA) + 8x 2.5- inch (NVMe) (HDD/SSD) maximum 184.32 TB Rear bays: Maximum 2x 2.5-inch SAS (SATA (NVMe) (HDD/S)
	Maximum 2x 2.5-inch SAS/SATA/NVMe (HDD/SSD) max 30.72 TB	• Maximum 2x 2.5-inch SAS/SATA/NVMe (HDD/ SSD) maximum 15.36 TB
Storage Controllers	Internal controllers: H965i, HBA465i (post RTS), H755, H755N, H355, HBA355i	Internal controllers: PERC H345, PERC H355, PERC H745, PERC H755, PERC H755N, HBA355i
	Internal Boot: Boot Optimized Storage Subsystem (BOSS N1): HWRAID 2x M.2 SSDs and Internal USB	Internal Boot: Internal Dual SD Module or Boot Optimized Storage Subsystem (BOSS S2): HWRAID 2x M.2 SSDs or Internal USB
	External: HBA355e, H965e and HBA465e (post	External PERC (RAID): PERC H840, HBA355e
	Software RAID: S160	Software RAID: S150



Table 2. Features comparison (continued)

Feature	PowerEdge R760xs	PowerEdge R750xs	
PCIe SSD	Front: Maximum 8 x 2.5-inch (NVMe drives)	Maximum 8 x 2.5-inch (NVMe drives)	
	Rear: Up to 2 x 2.5-inch NVMe		
PCIe Slots	Up to 6 PCIe slots (2 x Gen5, 4 x Gen4)	Up to 6 PCIe slots (5 x Gen4, 1 x Gen3)	
Embedded NIC (LOM)	2x 1GbE LOM	2x 1GbE LOM	
Networking Options (OCP 3.0)	Rear: 1 x OCP 3.0 (x8 PCle lanes)	Maximum 1 OCP 3.0 (x16 PCIe lanes)	
GPU	Nvidia A2 (60 W, LP)	Not supported	
I/O Ports	Front ports • 1x Dedicated iDRAC micro-USB • 1x USB 2.0 • 1x VGA	Front ports 1x Dedicated iDRAC micro-USB 1x USB 2.0 1x VGA 	
	Rear ports: • 1 x Dedicated iDRAC Ethernet port • 1x USB 2.0 • 1x USB 3.0 • 1x Serial (optional) • 1x VGA • 2x Ethernet	Rear ports: • 1 x Dedicated iDRAC Ethernet port • 1x USB 2.0 • 1x USB 3.0 • 1x Serial (optional) • 1x VGA • 2x Ethernet	
	Internal port: • 1x USB 3.0 (optional)	Internal port: • 1x USB 3.0 (optional)	
Rack Height	2U rack server	2U rack server	
Power Supplies	 1800 W Titanium 200–240 VAC/ 240 VDC 1400 W Platinum 100–240 VAC/ 240 VDC 1400 W Titanium 277 VAC/ 336 VDC 1100 W Titanium 100–240 VAC/ 240 VDC 1100 W DC/-48–(-60) V 800 W Platinum 100–240 VAC/ 240 VDC 700 W Titanium 200–240 VAC/ 240 VDC 600 W Platinum 100–240 VAC/ 240 VDC 	 1800 W Platinum 100–240 VAC/ 240 VDC 1400 W Platinum 100–240 VAC/ 240 VDC 1100 W Titanium 100–240 VAC/ 240 VDC 1100 W DC/-48–(-60) V 800 W Platinum 100–240 VAC/ 240 VDC 700 W Titanium 200–240 VAC/ 240 VDC 600 W Platinum 100–240 VAC/ 240 VDC 	
System Management	 Lifecycle Controller 3.x OpenManage QuickSync 2.0 OpenManage Enterprise Power Manager Digital License Key iDRAC Direct (dedicated micro-USB port) Easy Restore 	 Lifecycle Controller 3.x OpenManage QuickSync 2.0 OpenManage Enterprise Power Manager Digital License Key iDRAC Direct (dedicated micro-USB port) Easy Restore 	
Availability	Hot-plug drives	Hot-plug drives	
	Hot-plug redundant cooling	Hot-plug redundant cooling	
	Hot-plug redundant power supplies	Hot-plug redundant power supplies	
	BOSS-N1	IDSDM	
		BOSS S2	



Chassis views and features

Topics:

Chassis views



Front view of the system



Figure 1. Front view of 16 x 2.5-inch SAS/SATA + 8 x 2.5-inch NVMe drive system



Figure 2. Front view of 16 x 2.5-inch SAS/SATA drive system

Ŭ.

Figure 3. Front view of 8 x 2.5-inch SAS/SATA or NVMe drive system



Figure 4. Front view of 12 x 3.5-inch SAS/SATA drive system



1 11								
	顧問題言	-	國際觀		國際關係	C		ju I
110		16		. 6	國政部回		200 (B) (B) (B)	ū

Figure 5. Front view of 8 x 3.5-inch SAS/SATA drive system



Figure 6. Front view of no backplane configuration (0 drive system)

Left control panel view



Figure 7. Left control panel

Table 3. Left control panel

ltem	Indicator, button, or connector	lcon	Description
1	Status LED indicators	NA	Indicates the status of the system. For more information, see the Status LED indicators section.
2	System health and system ID	ž.	Indicates the system health. For more information, see the System health and system ID indicator codes section.





Figure 8. Left control panel with optional iDRAC Quick Sync 2 indicator

Table 4. Left control panel with optional iDRAC Quick Sync 2 indicator

ltem	Indicator, button, or connector	lcon	Description
1	Status LED indicators	N/A	Indicates the status of the system. For more information, see the Status LED indicators section.
2	System health and system ID indicator	ż	Indicates the system health. For more information, see the System health and system ID indicator codes section.
3	iDRAC Quick Sync 2 wireless indicator (optional)		Indicates if the iDRAC Quick Sync 2 wireless option is activated. The Quick Sync 2 feature allows management of the system using mobile devices. This feature aggregates hardware/ firmware inventory and various system level diagnostic/error information that can be used in troubleshooting the system. You can access system inventory, Dell Lifecycle Controller logs or system logs, system health status, and also configure iDRAC, BIOS, and networking parameters. You can also launch the virtual Keyboard, Video, and Mouse (KVM) viewer and virtual Kernel- based Virtual Machine (KVM), on a supported mobile device. For more information, see the Integrated Dell Remote Access Controller User's Guide at www.dell.com/poweredgemanuals.

(i) NOTE: For more information about the indicator codes, see the System diagnostics and indicator codes section.

Right control panel view



Figure 9. Right control panel



Table 5. Right control panel

ltem	Indicator or button	lcon	Description
1	Power button	Ċ	Indicates if the system is powered on or off. Press the power button to manually power on or off the system. Image: Note: Press the power button to gracefully shut down an ACPI-compliant operating system.
2	USB 2.0 port	4	The USB port is 4-pin, 2.0-compliant. This port enables you to connect USB devices to the system.
3	iDRAC Direct (Micro-AB USB) port	3.E	 The iDRAC Direct (Micro-AB USB) port enables you to access the iDRAC direct Micro-AB USB features. For more information, see the <i>Integrated Dell Remote Access Controller User's Guide</i> at www.dell.com/poweredgemanuals. NOTE: You can configure iDRAC Direct by using a USB to micro USB (type AB) cable, which you can connect to your laptop or tablet. Cable length should not exceed 3 feet (0.91 meters). Performance could be affected by cable quality.
4	VGA port	ici	Enables you to connect a display device to the system.

Rear view of the system

Figure 10. Rear view of the system



Figure 11. Rear view of the system with no riser and one CPU



Figure 12. Rear view of the system with no riser and two CPUs





Figure 13. Rear view of the system with Riser 1c



Figure 14. Rear view of the system with Riser 1d

Chassis views and features 13



Inside the system



Figure 15. Inside the system without rear drive and riser

- 1. Rear mounted front PERC
- 3. Memory module slots
- 5. Intrusion switch
- 7. PSU 1 and PSU 2
- 9. Processor heat sink
- 11. NVMe backplane

- 2. Cooling fan assembly
- 4. Power interposer board
- 6. OCP
- 8. System board
- 10. SAS/SATA backplane
- 12. Information tag





Figure 16. Inside the system with rear drive cage and riser

- 1. Rear mounted front PERC
- 3. Memory module slots
- 5. Intrusion switch
- 7. PSU 1 and PSU 2
- 9. System board
- 11. SAS/SATA backplane
- 13. Information tag

- 2. Cooling fan assembly
- 4. Power interposer board
- 6. Rear drive cage
- 8. Riser
- 10. Processor heat sink
- 12. NVMe backplane

System diagnostics and indicator codes

The diagnostic indicators on the system front panel display system status during system startup.

Status LED indicators

(i) NOTE: The indicators display solid amber if any error occurs.

Chassis views and features 15





Figure 17. Status LED indicators

Table 6.	Status	LED	indicators	and	descri	ptions
----------	--------	-----	------------	-----	--------	--------

lcon	Description	Condition	Corrective action		
0	Drive indicator	The indicator turns solid amber if there is a drive error.	 Check the System Event Log to determine if the drive has an error. Run the appropriate Online Diagnostics test. Restart the system and run embedded diagnostics (ePSA). If the drives are configured in a RAID array, restart the system, and enter the host adapter configuration utility program. 		
8	Temperature indicator	The indicator turns solid amber if the system experiences a thermal error (for example, the ambient temperature is out of range or there is a fan failure).	 Ensure that none of the following conditions exist: A cooling fan has been removed or has failed. System cover, air shrouds, or back filler bracket has been removed. Ambient temperature is too high. External airflow is obstructed. If the problem persists, see the Getting help section. 		
	Electrical indicator	The indicator turns solid amber if the system experiences an electrical error (for example, voltage out of range, or a failed power supply unit (PSU) or voltage regulator).	Check the System Event Log or system messages for the specific issue. If it is due to a problem with the PSU, check the LED on the PSU. Reseat the PSU. If the problem persists, see the Getting help section.		
ø	Memory indicator	The indicator turns solid amber if a memory error occurs.	Check the System Event Log or system messages for the location of the failed memory. Reseat the memory module. If the problem persists, see the Getting help section.		
IJ	PCIe indicator	The indicator turns solid amber if a PCIe card experiences an error.	Restart the system. Update any required drivers for the PCle card. Reinstall the card.If the problem persists, see the Getting help section.If the problem persists, see the Expansion cards and expansion card risers > Expansion card installation guidelines section.		



System health and system ID indicator codes

The system health and system ID indicator is located on the left control panel of the system.



Figure 18. System health and system ID indicator

Table 7. System health and system ID indicator codes

System health and system ID indicator code	Condition
Solid blue	Indicates that the system is powered on, is healthy, and system ID mode is not active. Press the system health and system ID button to switch to system ID mode.
Blinking blue	Indicates that the system ID mode is active. Press the system health and system ID button to switch to system health mode.
Solid amber	Indicates that the system is in fail-safe mode. If the problem persists, see the Getting help section.
Blinking amber	Indicates that the system is experiencing a fault. Check the System Event Log for specific error messages. For information about the event and error messages generated by the system firmware and agents that monitor system components, go to qrl.dell.com > Look Up > Error Code, type the error code, and then click Look it up.



QR code for PowerEdge R760xs system resources



Figure 19. QR code for PowerEdge R760xs system

18 Chassis views and features





Topics:

- Processor features
- Chipset

Processor features

The Intel 4th and 5th Generation Xeon[®] Scalable Processors stack is the next-generation data center processor offering with significant performance increases, integrated acceleration, and next-generation memory and I/O. Sapphire Rapids and Emerald Rapids accelerate customer usage with unique workload optimizations and provide the following feature improvements.

- Faster UPI with up to three Intel Ultra Path Interconnect (Intel UPI) at up to 20 GT/s, increasing multisocket bandwidth.
- More, faster I/O with PCI Express 5 and up to 80 lanes (per CPU)
- Enhanced Memory Performance with DDR5 support and memory speed up to 5200 MT/s in one DIMM per channel (1DPC).
- New onboard accelerators for data analytics, networking, storage, crypto, and data compression
- Enhanced security for virtualized environments with Intel Trust Domain Extensions (IntelR TDX) for confidential computing

Supported processors

The following table shows the Intel Sapphire Rapids and Emerald Rapids SKUs that are supported on the R760xs.

Processor	Clock Speed (GHz)	Cache (M)	UPI (GT/s)	Cores	Turbo	Memory Speed (MT/s)	TDP
6534	3.9	22.5	20	8	Turbo	4800	195 W
6526Y	2.8	37.5	20	16	Turbo	5200	195 W
5512U	2.1	52.5	N/A	28	Turbo	4800	185 W
4514Y	2.0	30	16	12	Turbo	4400	150 W
4510	2.4	30	16	12	Turbo	4400	150 W
4509Y	2.6	23	16	8	Turbo	4400	125 W
6448Y	2.2	60	16	32	Turbo	4800	225 W
6442Y	2.6	45	16	24	Turbo	4800	225 W
6438Y+	2	60	16	32	Turbo	4800	205 W
6426Y	2.6	30	16	16	Turbo	4800	185 W
6414U	2	60	16	32	Turbo	4800	250 W
5420+	2	53	16	28	Turbo	4400	205 W
5418Y	2	45	16	24	Turbo	4400	185 W
5416S	2	30	16	16	Turbo	4400	150 W
5415+	2.9	15	16	8	Turbo	4400	150 W
5412U	2.1	45	16	24	Turbo	4400	185 W
4416+	2	38	16	20	Turbo	4000	165 W

Table 8. Supported Processors for R760xs

Processor 19



Table 8. Supported Processors for R760xs (continued)

Processor	Clock Speed (GHz)	Cache (M)	UPI (GT/s)	Cores	Turbo	Memory Speed (MT/s)	TDP
4410Y	2	23	16	12	Turbo	4000	150 W
4410T	2.7	27	16	10	Turbo	4000	150 W
3408U	1.8	15	16	8	Turbo	4000	125 W

Chipset

The system supports Intel[®] C741 series chipset. DMI - 3.0 speed (port width x8, x4) USB ports - up to 10 superspeed (USB 3.1), 14 highspeed (USB 2.0)

SATA ports - up to 20 SATA port

PCIe Express - Up to 20 lanes, PCIe 3.0

Chipset features

- PCI-E interfaces
 - \circ $\;$ Integrated PCI Express Gen5 for improved bandwidth and connectivity
 - \circ $\,$ Up to 80 lanes per processor
 - Connect PCIe x1 to iDRAC- integrated VGA chip
- Integrated USB maximum of 10 SuperSpeed (USB 3.1), 14 highspeed (USB 2.0)
 - \circ $\,$ One front port (USB 2.0 / Right front I/O) $\,$
 - Two rear ports (USB 2.0/3.0)



Memory subsystem

Topics:

- Supported memory
- General memory module installation guidelines

Supported memory

Table 9. Memory technology comparison

Feature	PowerEdge R760xs (DDR5)
DIMM type	RDIMM
Transfer speed	5200 MT/s (1DPC) NOTE: Maximum DIMM transfer speed support depends on CPU SKU and DIMM population.
Voltage	1.1 V

Table 10. Supported memory matrix

DIMM type	Rank	Capacity	DIMM rated voltage	Operating Speed	
			and speed	1 DIMM per channel (DPC)	
RDIMM	1 R	16 GB	DDR5 (1.1 V), 4800 MT/s	Up to 4800 MT/s	
			DDR5 (1.1 V), 5600 MT/s	op to 5200 Min/s	
	2 R	32 GB, 64 GB, 96 GB	DDR5 (1.1 V), 4800 MT/s DDR5 (1.1 V), 5600 MT/s	Up to 4800 MT/s Up to 5200 MT/s	

(i) NOTE: 5600 MT/s RDIMMs are applicable for 5th Gen IntelR XeonR Scalable Processors.

(i) NOTE: The processor may reduce the performance of the rated DIMM speed.

General memory module installation guidelines

To ensure optimal performance of your system, observe the following general guidelines when configuring your system memory. If your system's memory configuration fails to observe these guidelines, your system might not boot, stop responding during memory configuration, or operate with reduced memory.

The memory bus may operate at speeds of 5200 MT/s, 4800 MT/s, 4400 MT/s or 4000 MT/s depending on the following factors:

- System profile selected (for example, Performance, Performance Per Watt Optimized (OS), or Custom [can be run at high speed or lower])
- Maximum supported DIMM speed of the processors



• Maximum supported speed of the DIMMs

(i) NOTE: MT/s indicates DIMM speed in MegaTransfers per second.

The system supports Flexible Memory Configuration, enabling the system to be configured and run in any valid chipset architectural configuration. The following are the recommended guidelines for installing memory modules:

- All DIMMs must be DDR5.
- Memory mixing is not supported for different DIMM capacities.
- If memory modules with different speeds are installed, they operate at the speed of the slowest installed memory module(s).
- Populate memory module sockets only if a processor is installed.
 - For single-processor systems, sockets A1 to A8 are available.
 - \circ $\,$ For dual-processor systems, sockets A1 to A8 and sockets B1 to B8 are available.
 - $\circ~$ A minimum of 1 DIMM must be populated for each installed processor.
- In **Optimizer Mode**, the DRAM controllers operate independently in the 64-bit mode and provide optimized memory performance.

Table 11. Memory population rules

Processor	Configuration	Memory population	Memory population information	
Single processor	Optimizer (Independent channel) population order	A{1}, A{2}, A{3}, A{4}, A{5}, A{6}, A{7}, A{8}	1, 2, 4, 6, 8 DIMMs are allowed.	
Dual processor (Start with processor1. Processor 1 and processor 2 population should match)		A{1}, B{1}, A{2}, B{2}, A{3}, B{3}, A{4}, B{4}, A{5}, B{5}, A{6}, B{6}, A{7}, B{7} A{8}, B{8}	 2, 4, 8, 12, 16 DIMMs are supported per system . i) NOTE: Optimizer population order is not traditional for 8 and 16 DIMMs installations for dual processor. 	

- Always populate memory channels identically with equal DIMMs for best performance.
- Supported RDIMM configurations are 1, 2, 4, 6, and 8 DIMMs per processor.
- Supported 96 GB RDIMM configurations are 1, 6 and 8 DIMMs per processor.
- Populate eight equal memory modules per processor (one DIMM per channel) at a time to maximize performance.

NOTE: Equal memory modules refer to DIMMs with identical electrical specification and capacity that may be from different vendors.




Topics:

- Storage controllers
- Supported Drives
- External Storage

Storage controllers

Dell RAID controller options offer performance improvements, including the fPERC solution. fPERC provides a base RAID HW controller without consuming a PCIe slot by using a small form factor and high-density connector to the base planar.

16G PERC Controller offerings are a heavy leverage of 15G PERC family. The Value and Value Performance levels carry over to 16G from 15G. New to 16G is the Avenger-based Premium Performance tier offering. This high-end offering drives IOPs performance and enhanced SSD performance.

Table 12. PERC Series controller offerings

Performance Level	Controller and Description
Entry	S160
Value	H355, HBA355e, HBA355i, (internal/external)
Value Performance	H755, H755N
Premium Performance	Н965і,
	Avenger 1
	Memory: 8GB DDR4 NV cache
	72-bit memory 2133 MHz
	Low profile form factors
	Dual A15 1.2 GHz CPU
	X8PCle 3.0, x8 12Gb SAS

() NOTE: For more information about the features of the Dell PowerEdge RAID controllers (PERC), Software RAID controllers, or BOSS card, and on deploying the cards, see the storage controller documentation at www.dell.com/storagecontrollermanuals.

(i) NOTE: From December 2021, H355 replaces H345 as the entry raid controller. H345 is deprecated in January 2022.

Storage controller feature matrix

Table 13. Storage controller feature matrix

Model & Form Factors	Interface Support	PCI Suppo rt	SAS Connection	Cach e Mem ory Size	Write Back Cache	RAID Levels	Max Drive Support	RAID Support	
	PowerEdge Server-Storage Controllers (PERC) Series 12								

Storage 23



Table 13. Storage controller feature matrix (continued)

Model & Form Factors	Interface Support	PCI Suppo rt	SAS Connection	Cach e Mem ory Size	Write Back Cache	RAID Levels	Max Drive Support	RAID Support
H965i Front	24Gb/s SAS 6Gb/s SAS/SATA Gen3 (8 GT/s) NVMe Gen4 (16 GT/s) NVMe	PCle Gen 4	16 ports/lanes - 2x8 Internal	8GB NV	Flash Backed Cache	0,1,5,6,10,50 ,60	16	Hardware
H965i Adapter	24Gb/s SAS 6Gb/s SAS/SATA Gen3 (8 GT/s) NVMe Gen4 (16 GT/s) NVMe	PCle Gen 4	16 ports/lanes - 2x8 Internal	8GB NV	Flash Backed Cache	0,1,5,6,10,50 ,60	16	Hardware
S160 Software RAID	Gen4 (16 GT/s) NVMe	PCle Gen 4	N/A	No Cach e	No Cache	0,1,5,10	8	Software RAID - Windows only
	PowerE	dge Serv	er-Storage Cont	rollers (PERC & SAS H	HBA) Series 11		
H755 Adapter	12Gb/s SAS 6Gb/s SAS/SATA 3Gb/s SAS/SATA Gen3 (8 GT/s) NVMe Gen4 (16 GT/s) NVMe	PCle Gen 4	16 ports- 2x8 Internal	8GB NV	Flash Backed Cache	0,1,5,6,10,50,60	16/ controller 50 with SAS Expander	Hardware
H755N Front (NVMe Only)	Gen3 (8 GT/s) NVMe Gen4 (16 GT/s) NVMe	PCle Gen 4	16 ports- 2x8 Internal	8GB NV	Flash Backed Cache	0,1,5,6,10,50 ,60	8/ controller	Hardware
H755 Front (SAS/ SATA only)	12Gb/s SAS 6Gb/s SAS/SATA 3Gb/s SAS/SATA	PCle Gen 4	16 ports- 2x8 Internal	8GB NV	Flash Backed Cache	0,1,5,6,10,50 ,60	16/ controller 50 with SAS Expander	Hardware
HBA355i Adapter	12Gb/s SAS 6Gb/s SAS/SATA 3Gb/s SAS/SATA	PCle Gen 4	16 ports- 2x8 Internal	N/A	N/A	N/A	16/ controller 50 with SAS Expander	N/A
HBA355i Front	12Gb/s SAS 6Gb/s SAS/SATA 3Gb/s SAS/SATA	PCle Gen 4	16 ports- 2x8 Internal	N/A	N/A	N/A	16/ controller 50 with SAS Expander	N/A

24 Storage

Inserido ao protocolo **22.951.206-4** por: **Pedro Henrique Golin Linhares** em: 31/03/2025 12:49. A autenticidade deste documento pode ser validada no endereço: https://www.eprotocolo.pr.gov.br/spiweb/validarDocumento com o código: **692a5ef82beb173ed33dad4ee6505546**.



Table 13. Storage controller feature matrix (continued)

Model & Form Factors	Interface Support	PCI Suppo rt	SAS Connection	Cach e Mem ory Size	Write Back Cache	RAID Levels	Max Drive Support	RAID Support
HBA355e Adapter	12Gb/s SAS 6Gb/s SAS/SATA 3Gb/s SAS/SATA	PCle Gen 4	16 ports- 4x4 external	NZA	N/A	N/A	240	N/A
H355 Adapter	12Gb/s SAS 6Gb/s SAS/SATA	PCle Gen 4	16 ports- 2x8 Internal	No Cach e	No Cache	0,1, 10	Up to 32 RAID, or 32 Non- RAID	Hardware
H355 Front	12Gb/s SAS 6Gb/s SAS/SATA	PCle Gen 4	16 ports- 2x8 Internal	No Cach e	No Cache	0,1, 10	Up to 32 RAID, or 32 Non- RAID	Hardware

() NOTE:

- 1. RAID 5/50 removed from entry RAID card
- 2. SWRAID support for Linus provides a pre-boot configuration utility to configure MDRAID and degraded boot capability.
- **3.** For information, post-RTS, see the Storage controller documentation at www.dell.com/stroagecontrollermanuals.

This document is updated as changes happen, so for the latest version be sure to bookmark it rather than downloading an offline copy or refer to the Storage Controller Matrix on sales portal.

Server storage controllers User Guide

• Server-Storage Controllers User's Guides, click here

RAID - Redundant Array of Independent Disks

• Link to Help Me Choose: RAID Configuration here

Datasheets and PERC performance scaling decks

- Resource Page for Server-Storage (Sales Portal) click here
- PERC & SAS HBA Datasheets (To be updated)

Boot Optimized Storage Solution (BOSS)

BOSS is a RAID solution that is designed to boot operating systems and segregate operating system boot drives from data on server-internal storage.



BOSS feature matrix

Table 14. BOSS feature matrix

BOSS card	Drive Size	RAID levels	Stripe size	Virtual disk cache functio n	Maxim um numbe r of virtual disks	Maxim um numbe r of drives suppor ted	Drive types	PCIe suppor t	Disk cache policy	Suppor t for Non- RAID disks	Crypto graphi c digital signatu re to verify firmwa re payloa d	Hot Plug
BOSS- N1 Monolit hic	M.2 devices are read- intensiv e with 480 GB or 960 GB capacit y	RAID1 and RAID0	Support s default 64K stripe size only	None	1	2	M.2 NVMe SSDs	Gen3	Drive default	No	Yes	Yes

BOSS-N1

BOSS-N1 is offered as a means of booting 16G servers to a full OS when the target OS is a full OS (not just a hypervisor), or the user does not wish to trade off standard hot plug drive slots for OS install.

The HW RAID BOSS-N1 card is a RAID controller with a limited feature set that presents M.2 NVMe-only SSDs as either a RAID0 disk or a single RAID1 volume with 2 disks. BOSS-N1 enables support for 480/960 GB Disks from Factory Install.

Hardware: BOSS-N1 Controller and Carrier (x2)

Reliability: Enterprise-Class M.2 NVMe SSDs

Supports dual 80 mm, Read Intensive (1DWPD), M.2 devices 480 GB/960 GB Standard - 1.92 TB QNS

Accessibility: Front Facing

Serviceability: Full Hot-Plug Support

Supports Hardware RAID1 and RAID0

Supports UEFI boot

Marvell 88NR2241 NVMe RAID Controller

Controlled Firmware Upgrade through iDRAC



Figure 20. BOSS-N1 Controller



Datasheets

• BOSS-N1 (to be updated)

BOSS User Guides

• BOSS-N1

Supported Drives

The table shown below lists the internal drives supported by the R760xs. Refer to Agile for the latest SDL

Table 15. Supported Drives

Form Factor	Туре	Speed	Rotational Speed	Capacities
2.5 inches	vSAS	12 Gb	SSD	1.92 TB, 3.84 TB, 960 GB, 7.62 TB
2.5 inches	SAS	24 Gb	SSD	1.92 TB, 1.6 TB, 800 GB, 3.84 TB, 960 GB, 7.68 TB
2.5 inches	SATA	6 Gb	SSD	1.92 TB, 480 GB, 960 GB, 3.84 TB
2.5 inches	NVMe	Gen4	SSD	1.6 TB, 3.2 TB, 6.4 TB, 1.92 TB, 3.84 TB, 15.63 TB, 7.68 TB
2.5 inches	DC NVMe	Gen4	SSD	3.84 TB, 960 GB
2.5 inches	SAS	12 Gb	10 K	600 GB, 1.2 TB, 2.4 TB
3.5 inches	SATA	6 Gb	7.2 K	2 TB, 4 TB, 8 TB, 12 TB, 16 TB, 20 TB
3.5 inches	SAS	12 Gb	7.2 K	2 TB, 4 TB, 8 TB, 12 TB, 16 TB, 20 TB

External Storage

The R760xs support the external storage device types listed in the table below.

Table 16. Support External Storage Devices

Device Type	Description			
External Tape	Supports connection to external USB tape products			
NAS/IDM appliance software	Supports NAS software stack			
JBOD	Supports connection to 12 Gb MD-series JBODs			





Topics:

- Overview
- OCP 3.0 support

Overview

PowerEdge offers a wide variety of options to get information moving to and from our servers. Industry best technologies are chosen, and systems management features are added by our partners to firmware to tie in with iDRAC. These adapters are rigorously validated for worry-free, fully supported use in Dell servers.

OCP 3.0 support

Table 17. OCP 3.0 feature list

Feature	OCP 3.0
Form factor	SFF
PCle Gen	Gen4
Max PCle width	x8
Max no. of ports	4
Port type	BT/SFP/SFP+/SFP28
Max port speed	25 GbE
NC-SI	Yes
SNAPI	No
WoL	Yes
Power consumption	15 W–35 W

Supported OCP cards

Table 18. Supported OCP cards

Form factor	Vendor	Port speed	Port type	Port count	DPN
OCP 3.0	Broadcom	10 GbE	BT	2	RN1M5
		25 GbE	SFP28	2	24FG6
		25 GbE	SFP28	4	3Y64D
		1 GbE	BT	4	VJWVJ
		10 GbE	ВТ	4	W5HC8
	Intel	1 GbE	ВТ	4	HY4CV

28 Networking



Table 18. Supported OCP cards (continued)

Form factor	Vendor	Port speed	Port type	Port count	DPN
		10 GbE	ВТ	2	F6X1R
		10 GbE	ВТ	4	XC0M4
		25 GbE	SFP28	2	PWH3C
		25 GbE	SFP28	4	Y4VV5

OCP NIC 3.0 vs. rack Network Daughter Card comparisons

Table 19. OCP 3.0, 2.0, and rNDC NIC comparison

Form Factor	Dell rNDC	OCP 2.0 (LOM Mezz)	OCP 3.0	Notes
PCle Gen	Gen 3	Gen 3	Gen 4	Supported OCP3 are SFF (small form factor)
Max PCle Lanes	x8	Up to x16	Up to x8	See server slot priority matrix
Shared LOM	Yes	Yes	Yes	This is iDRAC port redirect
Aux Power	Yes	Yes	Yes	Used for Shared LOM

OCP form factors



Figure 21. OCP 3.0 Small Card Form Factor (LS)

The process of installing the OCP card in R760xs system:

- 1. Open the blue latch on the system board.
- 2. Slide the OCP card into the slot in the system.
- **3.** Push until the OCP card is connected to the connector on the system board.
- **4.** Close the latch to lock the OCP card to the system.





Figure 22. Installing the OCP Card in R760xs

The process of removing the OCP card in R760xs system:

- 1. Open the blue latch to unlock the OCP card.
- 2. Push the OCP card towards the rear end of the system to disconnect from the connector on the system board.
- $\ensuremath{\textbf{3.}}$ Slide the OCP card out of the slot on the system.



Figure 23. Removing the OCP Card in R760xs



Slot priority matrix

For add-in cards that can be mapped to the R760xs and guidelines for installing expansion cards, see the R760xs slot priority matrix file on Sales Portal.

Link:https://www.delltechnologies.com/resources/en-us/auth/products/servers/category.htm

Topics:

• Expansion card installation guidelines

Expansion card installation guidelines



Figure 24. Expansion card slot connectors

- 1. SL9_CPU2_PB5 (PCIe cable connector for Riser 1C and Riser 1D)
- 2. SL10_CPU2_PA5 (PCIe cable connector for Riser 1C and Riser 1D)
- 3. SIG_PWR_0 (Power connector for Riser 1C and Riser 1D)
- 4. SL11_CPU1_PA6 (PCIe cable connector for Riser 1D)
- 5. SL12_CPU1_PB6 (PCIe cable connector for Riser 1D)

The following table describes the expansion card riser configurations:



Table 20. Expansion card riser configurations

Configuratio ns	Expansion card risers	PCIe Slots	Controlling processor	Height	Length	Slot width	Power
Config 0-1.	No riser	1, 2	Processor 1	Low profile	Half length	x16, x8	75 W
Config 0-2.	No riser	1, 2	Processor 1	Low profile	Half length	x16, x8	75 W
		5, 6	Processor 2	Low profile	Half length	x16, x16	75 W
Config 1.	R1C	1, 2, 3	Processor 1	Low profile	Half length	x16, x8, x16	75 W
		4, 5, 6	Processor 2	Low profile	Half length	x16, x16, x16	75 W
Config 2.	R1D	1, 2, 3, 4	Processor 1	Low profile	Half length	x16, x8, x8, x8	75 W

(i) NOTE: Only one cable riser can be installed at a time in any given configuration.

(i) NOTE: The slots 1, 2, 5 and 6 are Gen4 slots, slot 3 and 4 located on risers are Gen5 slots.



Figure 25. Riser 1C

- 1. Slot 3
- 2. Slot 4

32 Slot priority matrix





Figure 26. Riser 1D

- 1. Slot 3
- 2. Slot 4

(i) NOTE: The expansion-card slots are not hot-swappable.

The following table provides guidelines for installing expansion cards to ensure proper cooling and mechanical fit. The expansion cards with the highest priority should be installed first using the slot priority indicated. All the other expansion cards should be installed in the card priority and slot priority order.

Table 21. Configuration 0-1: No riser configuration

Card type	Slot priority	Maximum number of cards	
Dell Serial port module (LP)	2	1	
fPERC	Integrated slot	1	
InterN/AI PERC adapter	1	1	
Dell ExterN/Al Adapter	2, 1	2	
Mellanox (NIC: 400Gb)	Not supported	N/A	
Mellanox (NIC: 200Gb)	Not supported	N/A	
Mellanox (NIC: 100Gb)	1	1	
Mellanox HDR100 VPI	1	1	
Mellanox HDR VPI	1	1	
Broadcom (NIC: 100Gb)	1	1	
Intel (NIC: 100Gb)	1	1	
Broadcom (SFP: 25Gb)	2, 1	2	
Intel (NIC: 25Gb)	2, 1	2	
Qlogic (NIC: 25Gb)	Not supported	N/A	
Qlogic (NIC: 10Gb)	Not supported	N/A	
SolarFlare (NIC: 25Gb)	Not supported	N/A	
Broadcom (HBA: FC64)	2,1	2	
Broadcom (HBA: FC32)	2, 1	2	
Marvell (HBA: FC32)	2, 1	2	



Table 21. Configuration 0-1: No riser configuration (continued)

Card type	Slot priority	Maximum number of cards	
Emulex (HBA: FC32)	Not supported	N/A	
Avago (HBA: FC16)	Not supported	N/A	
Qlogic (HBA: FC16)	Not supported	N/A	
Broadcom (NIC: 10Gb)	2, 1	2	
Intel (NIC: 10Gb)	2, 1	2	
Qlogic (NIC: 10Gb)	Not supported	N/A	
Broadcom (NIC: 1Gb)	2, 1	N/A	
Intel (NIC: 1Gb)	2, 1	2	
Intel (OCP: 100Gb)	Integrated slot	1	
Broadcom (OCP: 25Gb)	Integrated slot	1	
Intel (OCP: 25Gb)	Integrated slot	1	
Marvell (OCP: 25Gb)	Not supported	NZA	
SolarFlare (OCP: 25Gb)	Not supported	N/A	
Broadcom (OCP: 10Gb)	Integrated slot	1	
Marvell (OCP: 10Gb)	Not supported	N/A	
Intel (OCP: 10Gb)	Not supported	N/A	
Broadcom (OCP: 1Gb)	Integrated slot	1	
Intel (OCP: 1Gb)	Not supported	N/A	
Dell BOSS N1 Module	Integrated slot	1	

Table 22. Configuration 0-2: No riser configuration

Card type	Slot priority	Maximum number of cards			
Dell Serial port module (LP)	2	1			
fPERC	Integrated slot	1			
InterN/AI PERC adapter	1	1			
Dell ExterN/Al Adapter	6, 2, 1, 5	4			
Mellanox (NIC: 400Gb)	Not supported	N/A			
Mellanox (NIC: 200Gb)	Not supported	N/A			
Mellanox (NIC: 100Gb)	6, 1, 5	3			
Mellanox HDR100 VPI	6, 1, 5	3			
Mellanox HDR VPI	6, 1, 5	3			
Broadcom (NIC: 100Gb)	6, 1, 5	3			
Intel (NIC: 100Gb)	6, 1, 5	3			
Broadcom (SFP: 25Gb)	6, 2, 1, 5	4			
Intel (NIC: 25Gb)	2, 1	2			
Qlogic (NIC: 25Gb)	Not supported	N/A			
Qlogic (NIC: 10Gb)	Not supported	N/A			
SolarFlare (NIC: 25Gb)	Not supported	N/A			

34 Slot priority matrix

Inserido ao protocolo **22.951.206-4** por: **Pedro Henrique Golin Linhares** em: 31/03/2025 12:49. A autenticidade deste documento pode ser validada no endereço: https://www.eprotocolo.pr.gov.br/spiweb/validarDocumento com o código: **692a5ef82beb173ed33dad4ee6505546**.



Table 22. Configuration 0-2: No riser configuration (continued)

Card type	Slot priority	Maximum number of cards	
Broadcom (HBA: FC64)	6, 2, 1, 5	4	
Broadcom (HBA: FC32)	6, 2, 1, 5	4	
Marvell (HBA: FC32)	6, 2, 1, 5	4	
Emulex (HBA: FC32)	Not supported	N/A	
Avago (HBA: FC16)	Not supported	N/A	
Qlogic (HBA: FC16)	Not supported	N/A	
Broadcom (NIC: 10Gb)	6, 2, 1, 5	4	
Intel (NIC: 10Gb)	6, 2, 1, 5	4	
Qlogic (NIC: 10Gb)	Not supported	N/A	
Broadcom (NIC: 1Gb)	6, 2, 1, 5	4	
Intel (NIC: 1Gb)	6, 2, 1, 5	4	
Intel (OCP: 100Gb)	Integrated slot	1	
Broadcom (OCP: 25Gb)	Integrated slot	1	
Intel (OCP: 25Gb)	Integrated slot	1	
Marvell (OCP: 25Gb)	Not supported	N/A	
SolarFlare (OCP: 25Gb)	Not supported	N/A	
Broadcom (OCP: 10Gb)	Integrated slot	1	
Marvell (OCP: 10Gb)	Not supported	N/A	
Intel (OCP: 10Gb)	Not supported	N/A	
Broadcom (OCP: 1Gb)	Integrated slot	1	
Intel (OCP: 1Gb)	Not supported	N/A	
Dell BOSS N1 Module	Integrated slot	1	

Table 23. Configuration 1: R1C

Card type	Slot priority	Maximum number of cards	
Dell Serial port module (LP)	2	1	
InterN/AI PERC adapter	1	1	
Dell exterN/AI PERC adapter	6, 2, 1, 3, 5, 4	6	
12Gbps SAS HBA	1	1	
Mellanox (NIC: 400Gb)	4, 3	2	
Mellanox (NIC: 200Gb)	4, 3	2	
Broadcom (NIC: 100Gb)	6, 1, 3, 5, 4	5	
Intel (NIC: 100Gb)	6, 1, 3, 5, 4	5	
Mellanox (NIC: 100Gb)	6, 1, 3, 5, 4	5	
Mellanox HDR100 VPI	6, 1, 3, 5, 4	5	
Mellanox HDR VPI	6, 1, 3, 5, 4	5	
Intel (NIC: 25Gb)	6, 2, 1, 3, 5, 4	6	
Mellanox (NIC: 25Gb)	6, 1, 3, 5, 4	5	



Table 23. Configuration 1: R1C (continued)

Card type	Slot priority	Maximum number of cards	
Qlogic (NIC: 25Gb)	Not Supported	N/A	
Broadcom (NIC: 10Gb)	6, 2, 1, 3, 5, 4	6	
Broadcom (NIC: 25Gb)	6, 2, 1, 3, 5, 4	6	
SolarFlare (NIC: 25Gb)	Not Supported	N/A	
Broadcom (HBA: FC64)	6, 2, 1, 3, 5, 4	6	
Broadcom (HBA: FC32)	6, 2, 1, 3, 5, 4	6	
QLogic (HBA: FC32)	6, 2, 1, 3, 5, 4	6	
Marvell (HBA: FC32)	6, 2, 1, 3, 5, 4	6	
Emulex (HBA: FC32)	Not Supported	N/A	
Avago (HBA: FC16)	Not Supported	N/A	
QLogic (HBA: FC16)	Not Supported	N/A	
Intel (NIC: 10Gb)	6, 2, 1, 3, 5, 4	6	
Qlogic (NIC: 10Gb)	Not Supported	N/A	
Intel (NIC: 1Gb)	6, 2, 1, 3, 5, 4	6	
Intel (OCP: 100Gb)	Integrated slot	1	
Broadcom (OCP: 25Gb)	Integrated slot	1	
Intel (OCP: 25Gb)	Integrated slot 1		
Marvell (OCP: 25Gb)	Not supported	N/A	
SolarFlare (OCP: 25Gb)	Not supported	N/A	
Broadcom (OCP: 10Gb)	Integrated slot	1	
Marvell (OCP: 10Gb)	Not supported	N/A	
Intel (OCP: 10Gb)	Not supported	N/A	
Broadcom (OCP: 1Gb)	Integrated slot	1	
Intel (OCP: 1Gb)	Not supported	N/A	
Dell BOSS N1 Module	Integrated slot	1	

Table 24. Configuration 2: R1D

Card type	Slot priority	Maximum number of cards		
Dell Serial port module (LP)	2	1		
InterN/AI PERC adapter	1	1		
Dell exterN/Al PERC adapter	4, 3, 2, 1	 		
12Gbps SAS HBA	1	1		
Mellanox (NIC: 400Gb)	Not supported	N/A		
Mellanox (NIC: 200Gb)	Not supported	N/A		
Broadcom (NIC: 100Gb)	1	1		
Intel (NIC: 100Gb)	1	1		
Mellanox (NIC: 100Gb)	1	1		
Mellanox HDR100 VPI	1	1		



Table 24. Configuration 2: R1D (continued)

Card type	Slot priority	Maximum number of cards	
Mellanox HDR VPI	1	1	
Intel (NIC: 25Gb)	4, 3, 2, 1	4	
Mellanox (NIC: 25Gb)	4, 3, 2, 1	4	
Qlogic (NIC: 25Gb)	Not Supported	N/A	
Broadcom (NIC: 10Gb)	4, 3, 2, 1	4	
Broadcom (NIC: 25Gb)	4, 3, 2, 1	4	
SolarFlare (NIC: 25Gb)	Not Supported	N/A	
Broadcom (HBA: FC64)	4, 3, 2, 1	4	
Broadcom (HBA: FC32)	4, 3, 2, 1	4	
QLogic (HBA: FC32)	4, 3, 2, 1	4	
Marvell (HBA: FC32)	4, 3, 2, 1	4	
Emulex (HBA: FC32)	Not Supported	N/A	
Avago (HBA: FC16)	Not Supported	N/A	
QLogic (HBA: FC16)	Not Supported	N/A	
Intel (NIC: 10Gb)	4, 3, 2, 1	4	
Qlogic (NIC: 10Gb)	Not Supported	N/A	
Intel (NIC: 1Gb)	4, 3, 2, 1	4	
Intel (OCP: 100Gb)	Integrated slot	1	
Broadcom (OCP: 25Gb)	Integrated slot	1	
Intel (OCP: 25Gb)	Integrated slot	1	
Marvell (OCP: 25Gb)	Not supported	N/A	
SolarFlare (OCP: 25Gb)	Not supported	N/A	
Broadcom (OCP: 10Gb)	Integrated slot	1	
Marvell (OCP: 10Gb)	Not supported	N/A	
Intel (OCP: 10Gb)	Not supported	N/A	
Broadcom (OCP: 1Gb)	Integrated slot	1	
Intel (OCP: 1Gb)	Not supported	N/A	
Dell BOSS N1 Module	Integrated slot	1	

(i) NOTE: The serial COM card is not a real PCIe add-in card and has a dedicated slot on the system board.



Power, thermal, and acoustics

PowerEdge servers have an extensive collection of sensors that automatically track thermal activity, which helps to regulate temperature by reducing server noise and power consumption. The table below lists the tools and technologies Dell offers to lower power consumption and increase energy efficiency.

Topics:

- Power
- Thermal
- Acoustics

Power

Table 25. Power tools and technologies

Feature	Description
Power Supply Units(PSU) portfolio	Dell's PSU portfolio includes intelligent features such as dynamically optimizing efficiency while maintaining availability and redundancy. Find additional information in the Power supply units section.
Tools for right sizing	Enterprise Infrastructure Planning Tool (EIPT) is a tool that can help you determine the most efficient configuration possible. With Dell's EIPT, you can calculate the power consumption of your hardware, power infrastructure, and storage at a given workload. Learn more at www.dell.com/calc.
Industry Compliance	Dell's servers are compliant with all relevant industry certifications and guide lines, including 80 PLUS, Climate Savers and ENERGY STAR.
Power monitoring accuracy	PSU power monitoring improvements include:
	 Dell's power monitoring accuracy is currently 1%, whereas the industry standard is 5% More accurate reporting of power Better performance under a power cap
Power capping	Use Dell's systems management to set the power cap limit for your systems to limit the output of a PSU and reduce system power consumption. Dell is the first hardware vendor to leverage Intel Node Manager for circuit-breaker fast capping.
Systems Management	iDRAC Enterprise and Datacenter provides server-level management that monitors, reports and controls power consumption at the processor, memory and system level.
	Dell OpenManage Power Center delivers group power management at the rack, row, and data center level for servers, power distribution units, and uninterruptible power supplies.
Active power management	Intel Node Manager is an embedded technology that provides individual server-level power reporting and power limiting functionality. Dell offers a complete power management solution comprised of Intel Node Manager accessed through Dell iDRAC9 Datacenter and OpenManage Power Center that allows policy-based management of power and thermal at the individual server, rack, and data center level. Hot spare reduces power consumption of redundant power supplies. Thermal control off a speed optimizes the thermal settings for your environment to reduce fan consumption and lower system power consumption. Idle power enables Dell servers to run as efficiently when idle as when at full workload.
Rack infrastructure	Dell offers some of the industry's highest-efficiency power infrastructure solutions, including:



Table 25. Power tools and technologies (continued)

Feature	Description
	 Power distribution units (PDUs) Uninterruptible power supplies (UPSs) Energy Smart containment rack enclosures Find additional information at: https://www.delltechnologies.com/en-us/servers/power-and-cooling.htm.

PSU specifications

The PowerEdge R760xs system supports up to two AC or DC power supply units (PSUs).

Table 26. R760xs PSU specifications	Table	26.	R760xs	PSU s	specifications	
-------------------------------------	-------	-----	---------------	-------	----------------	--

PSU Clas		Heat dissipation	Frequ ency	AC Voltag	e	DC Voltage	-48 V DC	277 V AC and	HVDC
		(maximum) (BTU/ hr)	(HZ)	Low Line AC (100– 120 V)	High Line AC (200– 240 V)	240 V DC	-40 V to -72 V DC	277 V AC (249 V AC– 305 V AC)	336 V (260 V DC-400 V DC)
600 W mixed mode	Platinu m	2250	50/60	600 W	600 W	600 W	N/A	N/A	N/A
700 W mixed mode HLAC	Titaniu m	2625	50/60	N/A	700 W	700 W	N/A	NZA	N/A
800 W mixed mode	Platinu m	3000	50/60	800 W	800 W	800 W	N/A	N/A	N/A
1100 W -48 V DC	N/A	4265	N/A	N/A	N/A	N/A	1100 W	N/A	N/A
1100 W mixed mode	Titaniu m	4125	50/60	1050 W	1100 W	1100 W	N/A	N/A	N/A
1400 W mixed mode	Titaniu m	5250	50/60	1050 W	1400 W	1400 W	N/A	N/A	N/A
1400 W mixed mode	Platinu m	5250	50/60	1050 W	1400 W	1400 W	N/A	N/A	N/A
1400 W 277 V AC and HVDC	Titaniu m	5250	50/60	N/A	N/A	N/A	N/A	1400 W	1400 W
1800 W mixed mode HLAC	Titaniu m	6610	50/60	N/A	1800 W	1800 W	N/A	N/A	N/A

(i) NOTE: Heat dissipation is calculated using the PSU wattage rating.

(i) NOTE: HLAC stands for High-Line AC, with a range of 200 - 240V AC. HVDC stands for High-Voltage DC, with 336 V DC.

NOTE: When selecting or upgrading the system configuration, to ensure optimum power utilization, verify the system power consumption with the Enterprise Infrastructure Planning Tool available at Dell.com/calc.

NOTE: If a system with AC 1400 W or 1100 W PSUs operates at low line 100-120 Vac, and then the power rating per PSU is degraded to 1050 W.





Figure 27. PSU power cords



Figure 28. APP 2006G1 power cord



Figure 29. Lotes DC PSU connector

Table 27. PSU power cords

Form factor	Output	Power cord
Redundant 60 mm	600 W Mixed Mode	C13
	700 W Mixed Mode HLAC	C13
	800 W Mixed Mode	C13
	1100 W Mixed Mode	C13
	1100 W -48 V DC	Lotes DC PSU connector
	1400 W Mixed Mode	C13
	1400 W 277 VAC and 336 VDC	APP 2006G1
	1800 W Mixed Mode HLAC	C15

(i) NOTE: C13 power cord combined with C14 to C15 jumper power cord can be used to adapt 1800 W PSU.



Thermal

PowerEdge servers have an extensive collection of sensors that automatically track thermal activity, which helps regulate temperature thereby reducing server noise and power consumption.

Acoustics

Acoustical configurations of R760xs

Dell PowerEdge R760xs is a rack server appropriate for attended data center environment. However, lower acoustical output is attainable with proper hardware or software configurations.

Configuration	Quietest	12 x 3.5-inch configuration	Quietest - GPU	Volume 1 - 2.5- inch	Volume 2 - 3.5- inch with Rear Storage
CPU TDP	125 W	185 W	150 W	185 W	185 W
CPU Quantity	1	2	2	2	2
RDIMM Memory	16 GB DDR5	32 GB DDR5	16 GB DDR5	16 GB DDR5	32 GB DDR5
Memory Quantity	2	16	4	16	16
Backplane Type	8 x 2.5-inch BP	12 x 3.5-inch BP	8 x 2.5-inch BP	8 x 2.5-inch + 8 x 2.5-inch NVMe BP	12 x 3.5-inch + rear 2 x 2.5-inch BP
HDD Type	SATA 2.5-inch 480 GB	SATA 3.5-inch 12 TB	SATA 2.5-inch 480 GB	SAS 2.5-inch + NVMe 2.5-inch	SATA 3.5-inch 12 TB + rear 2.5-inch U.2 NVMe
HDD Quantity	1	12	4	8+8	12+2
Flash Drives	×	×	X	×	X
Flash Quantity	×	×	X	×	X
PSU Type	600 W	1400 W	800 W	800 W	1400 W
PSU Quantity	1	2	2	2	2
OCP	1G	10/25 2-port	1G	10/25 2-port	10/25 2-port
PCI 1	×	×	X	×	X
PCI 2	X	2-port 25Gb	X	2-port 25Gb	2-port 25Gb
PCI 3	X	×	X	X	X
PCI 4	X	X	A2	X	×
PCI 5	X	2-port 25 Gb	X	2-port 25 Gb	2-port 25 Gb
PCI 6	X	×	×	X	X
PERC	Front H355i	Front H755	Front H355i	Front H755	Rear H755

Table 28. Configurations tested for acoustical experience

Table 29. Acoustical experience of R760 configurations

Configuratio	n	Quietest	12 x 3.5-inch configuration	Quietest - GPU	Volume 1 - 2.5- inch	Volume 2 - 3.5-inch with Rear Storage
Acoustical Performance: Idle/ Operating @ 25°C Ambient						
L _{wA,m} (B)	Idle ⁽⁴⁾	4.6	7.0	6.4	6.4	7.2



Table 29. Acoustical experience of R760 configurations (continued)

Configuratio	n	Quietest	12 x 3.5-inch configuration	Quietest - GPU	Volume 1 - 2.5- inch	Volume 2 - 3.5-inch with Rear Storage
	Operating/Customer usage operating ⁽⁵⁾⁽⁶⁾	4.6	7.0	8.0	6.4	7.2
К _v (В)	Idle (4)	0.4	0.4	0.4	0.4	0.4
	Operating/Customer usage operating ⁽⁵⁾⁽⁶⁾	0.4	0.4	0.4	0.4	0.4
L _{pA,m} (dB)	Idle ⁽⁴⁾	32	56	50	49	57
	Operating/Customer usage operating ⁽⁵⁾⁽⁶⁾	32	56	64	50	58
Prominent dis	crete tones ⁽³⁾	Prominence ratio < ECMA-74	e Prominence ratio < 15 dB			
Acoustical Pe	rformance: Idle @ 28°C /	Ambient	•			
L _{wA,m} ⁽¹⁾ (B)		4.8	7.2	6.6	6.6	7.4
К _v (В)		0.4	0.4	0.4	0.4	0.4
L _{pA,m} ⁽²⁾ (dB)		33	56	52	52	58
Acoustical Performance: Max. loading @ 35°C Ambient						
L _{wA,m} ⁽¹⁾ (B)		5.4	8.0	8.6	8.1	8.1
К _v (В)		0.4	0.4	0.4	0.4	0.4
$L_{pA,m}^{(2)}(dB)$		38	66	70	65	65

⁽¹⁾LwA,m: The declared mean A-weighted sound power level (LwA) is calculated per section 5.2 of ISO 9296 with data collected using the methods described in ISO 7779 (2010). Engineering data presented here may not be fully compliant with ISO 7779 declaration requirements.

⁽²⁾LpA,m: The declared mean A-weighted emission sound pressure level is at the bystander position per section 5.3 of ISO 9296 and measured using methods described in ISO 7779. The system is placed in a 24U rack enclosure, 25cm above a reflective floor. Engineering data presented here may not be fully compliant with ISO 7779 declaration requirements.

⁽³⁾Prominent tones: Criteria of Annex D of ECMA-74 & Prominence Ratio method of ECMA-418 are followed to determine if discrete tones are prominent and to report them, if so.

⁽⁴⁾Idle mode: Idle mode is the steady-state condition in which the server is energized but not operating any intended function.

⁽⁵⁾Operating mode: Operating mode is represented by the maximum of the steady state acoustical output at 50% of CPU TDP or active storage drives for the respective sections of Annex C of ECMA-74.

⁽⁶⁾ Customer Usage Operating mode: The operating mode is represented by the maximum of the steady state acoustical output at 10%~25% of CPU TDP, 0%~10% IOPs load, and >80% GPU load as the components showed in the above configurations.



Rack, rails, and cable management

Topics:

- Rails information
- Cable Management Arm
- Strain Relief Bar
- Rack Installation

Rails information

The rail offerings for the R760xs consist of two general types: sliding and static. The cable management offerings consist of an optional cable management arm (CMA) and an strain relief bar (SRB).

See the Dell Enterprise Systems Rail Sizing and Rack Compatibility Matrix available at rail-rack-matrix for information regarding:

- Specific details about rail types.
- Rail adjustability ranges for various rack mounting flange types
- Rail depth with and without cable management accessories
- Rack types that are supported for various rack mounting flange types

Key factors governing proper rail selection include the following:

- Identifying the type of rack in which they will be installed.
- The spacing between the front and rear mounting flanges of the rack.
- The type and location of any equipment mounted in the back of the rack such as power distribution units (PDUs), and the overall depth of the rack
- Overall depth of the rack

A11 Sliding Rails features summary

The sliding rails allow the system to be fully extended out of the rack for service. The sliding rails have a Cable Management Arm (CMA) and a Strain Relief Bar (SRB) option.

There are one types of sliding rails available:

• Stab-in/Drop-in sliding rails

A11 Stab-in/Drop-in sliding rails for 4-post racks

- Supports Drop-in or Stab-in installation of the chassis to the rails
- Supports tool-less installation in 19-inch EIA-310-E compliant square, unthreaded round hole racks including all generations of Dell racks.

Also supports tool-less installation in threaded round hole 4-post racks

- Support for tool-less installation in Dell Titan or Titan-D racks
- Supports full extension of the system out of the rack to allow serviceability of key internal components
- (i) **NOTE:** For situations where CMA support is not required, the outer CMA mounting brackets can be uninstalled from the sliding rails. This reduces the overall length of the rails and eliminates the potential interferences with rear mounted PDUs or the rear rack door.

Supports optional Cable Management Arm (CMA)

• Supports optional Strain Relief Bar (SRB)



A8 Static Rails features summary

The static rails, which are shown in the figure below, support a wider variety of racks than the sliding rails, but do not support serviceability in the rack. The static rails are not compatible with the CMA and SRB.

- Supports Stab-in installation of the chassis to the rails
- Supports tool-less installation in 19-inch EIA-310-E compliant square or unthreaded round hole 4-post racks including all generations of Dell racks
- Supports tooled installation in 19-inch EIA-310-E compliant threaded hole 4-post and 2-post racks
- Supports tooled installation in Dell Titan or Titan-D rack

() NOTE:

- Screws are not included with the static rail kit since racks are offered with various thread types.
- Screw head diameter should be 10mm or less.



Figure 30. Static rails

2-Post racks installation

If installing to 2-Post (Telco) racks, the ReadyRails II static rails (A8) must be used. Sliding rails support mounting in 4-post racks only.



Figure 31. Static rails in 2-post center mount configuration

44 Rack, rails, and cable management



Installation in the Dell Titan or Titan-D racks

For tool-less installation in Titan or Titan-D racks, the Stab-in/Drop-in sliding rails (A11) must be used. This rail collapses down sufficiently to fit in the rack with mounting flanges that are spaced about 24 inches apart from front to back. The Stab-in/Drop-in sliding rail allows bezels of the servers and storage systems to be aligned when installed in these racks. For tooled installation, Stab-in Static rails (A8) must be used for bezel alignment with storage systems.

Cable Management Arm

The optional Cable Management Arm (CMA) for the system organizes and secures the cords and cables exiting the back of the server and unfolds to allow the server to extend out of the rack without having to detach the cables.

Some key features of the CMA include:

- Large U-shaped baskets to support dense cable loads
- Open vent pattern for optimal airflow
- Support for mounting on either side by swinging the spring-loaded brackets from one side to the other
- Utilizes hook-and-loop straps rather than plastic tie wraps to eliminate the risk of cable damage during cycling
- Includes a low-profile fixed tray to both support and retain the CMA in its fully closed position
- Both the CMA and the tray mount without the use of tools through simple and intuitive snap-in designs

The CMA can be mounted to either side of the sliding rails without the use of tools or the need for conversion. For systems with one power supply unit (PSU), it is recommended to mount on the side opposite to that of the power supply to allow easier access to it and the rear drives (if applicable) for service or replacement.



Figure 32. Cable Management Arm

Strain Relief Bar

The optional strain relief bar (SRB) for the R760xs organizes and supports cable connections at the rear end of the server to avoid damage from bending.





Figure 33. Cabled strain relief bar

Sliding rails with optional SRB:

- Support tool-less attachment to rails
- Support two depth positions to accommodate various cable loads and rack depths
- Support cable loads and controls stress on server connectors
- Support cables can be segregated into discrete, purpose-specific bundles

Rack Installation

Drop-in design means that the system is installed vertically into the rails by inserting the standoffs on the sides of the system into the J-slots in the inner rail members with the rails in the fully extended position. The recommended method of installation is to first insert the rear standoffs on the system into the rear J-slots on the rails to free up a hand and then rotate the system down into the remaining J-slots while using the free hand to hold the rail against the side of the system.

Stab-in design means that the inner (chassis) rail members must first be attached to the sides of the system and then inserted into the outer (cabinet) members installed in the rack.

Installing system into the rack (option A: Drop-In)

1. Pull the inner rails out of the rack until they lock into place.





Figure 34. Pull out inner rail

- 2. Locate the rear rail standoff on each side of the system and lower them into the rear J-slots on the slide assemblies.
- 3. Rotate the system downward until all the rail standoffs are seated in the J-slots.



Figure 35. Rail standoffs seated in J-slots

- 4. Push the system inward until the lock levers click into place.
- 5. Pull the blue side release lock tabs forward or backward on both rails and slide the system into the rack until the system is in the rack.





Figure 36. Slide system into the rack

Installing the system into the rack (option B: Stab-In)

- 1. Pull the intermediate rails out of the rack until they lock into place.
- 2. Release the inner rail lock by pulling forward on the white tabs and sliding the inner rail out of the intermediate rails.



Figure 37. Pull out the intermediate rail

Table 30. Rail component label

Number	Component
1	Intermediate rail
2	Inner rail

3. Attach the inner rails to the sides of the system by aligning the J-slots on the rail with the standoffs on the system and sliding forward on the system until they lock into place.





Figure 38. Attach the inner rails to the system

4. With the intermediate rails extended, install the system into the extended rails.



Figure 39. Install system into the extended rails

5. Pull blue slide release lock tabs forward or backward on both rails, and slide the system into the rack.





Figure 40. Slide system into the rack

50 Rack, rails, and cable management

Inserido ao protocolo **22.951.206-4** por: **Pedro Henrique Golin Linhares** em: 31/03/2025 12:49. A autenticidade deste documento pode ser validada no endereço: https://www.eprotocolo.pr.gov.br/spiweb/validarDocumento com o código: **692a5ef82beb173ed33dad4ee6505546**.



Supported operating systems

The PowerEdge R760xs system supports the following operating systems:

- Canonical Ubuntu Server LTS
- Microsoft Windows Server with Hyper-V
- Red Hat Enterprise Linux
- SUSE Linux Enterprise Server
- VMware vSAN/ESXi

For more information, go to www.dell.com/ossupport.

Supported operating systems 51



Dell OpenManage Systems Management

Dell delivers management solutions that help IT administrators effectively deploy, update, monitor, and manage IT assets. OpenManage solutions and tools enable you to quickly respond to problems by helping them to manage Dell servers efficiently; in physical, virtual, local, and remote environments; all without the need to install an agent in the operating system.

The OpenManage portfolio includes:

- Innovative embedded management tools integrated Dell Remote Access Controller (iDRAC)
- Consoles OpenManage Enterprise
- Extensible with plug-ins OpenManage Power Manager
- Update tools Repository Manager

Dell has developed comprehensive systems management solutions that are based on open standards and has integrated with management consoles from partners such as Microsoft and VMware, allowing advanced management of Dell servers. Dell management capabilities extend to offerings from the industry's top systems management vendors and frameworks such as Ansible, Splunk, and ServiceNow. OpenManage tools automate the full span of server life cycle management activities along with powerful RESTful APIs to script or integrate with your choice of frameworks.

For more information about the entire OpenManage portfolio, see:

• The latest Dell Systems Management Overview Guide.

Topics:

- Integrated Dell Remote Access Controller (iDRAC)
- Systems Management software support matrix

Integrated Dell Remote Access Controller (iDRAC)

iDRAC9 delivers advanced, agent-free, local and remote server administration. Embedded in every PowerEdge server, iDRAC9 provides a secure means to automate a multitude of common management tasks. Because iDRAC is embedded within every PowerEdge server, there is no additional software to install; just plug in power and network cables, and iDRAC is ready to go. Even before installing an operating system (operating system) or hypervisor, IT administrators have a complete set of server management features at their fingertips.

With iDRAC9 in-place across the Dell PowerEdge portfolio, the same IT administration techniques and tools can be applied throughout. This consistent management platform allows easy scaling of PowerEdge servers as an organization's infrastructure grows. Customers can use the iDRAC RESTful API for the latest in scalable administration methods of PowerEdge servers. With this API, iDRAC enables support for the Redfish standard and enhances it with Dell extensions to optimize at-scale management of PowerEdge servers. By having iDRAC at the core, the entire OpenManage portfolio of Systems Management tools allows every customer to tailor an effective, affordable solution for any size environment.

Zero Touch Provisioning (ZTP) is embedded in iDRAC. ZTP - Zero Touch Provisioning is Intelligent Automation Dell's agent-free management puts IT administrators in control. Once a PowerEdge server is connected to power and networking, that system can be monitored and fully managed, whether you're standing in front of the server or remotely over a network. In fact, with no need for software agents, an IT administrator can: • Monitor • Manage • Update • Troubleshoot and remediate Dell servers With features like zero-touch deployment and provisioning, iDRAC Group Manager, and System Lockdown, iDRAC9 is purpose-built to make server administration quick and easy. For those customers whose existing management platform utilizes in-band management, Dell does provide iDRAC Service Module, a lightweight service that can interact with both iDRAC9 and the host operating system to support legacy management platforms.

When ordered with DHCP enabled from the factory, PowerEdge servers can be automatically configured when they are initially powered up and connected to your network. This process uses profile-based configurations that ensure each server is configured per your specifications. This feature requires an iDRAC Enterprise license.

iDRAC9 offers following license tiers:



Table 31. iDRAC9 license tiers

License	Description
iDRAC9 Basic	 Available only on 100-500 series rack/tower Basic instrumentation with iDRAC web UI For cost conscious customers that see limited value in management
iDRAC9 Express	 Default on 600+ series rack/tower, modular, and XR series Includes all features of Basic Expanded remote management and server life-cycle features
iDRAC9 Enterprise	 Available as an upsell on all servers Includes all features of Basic and Express. Includes key features such as virtual console, AD/LDAP support, and more Remote presence features with advanced, Enterprise-class, management capabilities
iDRAC9 Datacenter	 Available as an upsell on all servers Includes all features of Basic, Express, and Enterprise. Includes key features such as telemetry streaming, Thermal Manage, automated certificate management, and more Extended remote insight into server details, focused on high end server options, granular power, and thermal management

For a full list of iDRAC features by license tier, see Integrated Dell Remote Access Controller 9 User's Guide at Dell.com. For more details on iDRAC9 including white papers and videos, see:

• Support for Integrated Dell Remote Access Controller 9 (iDRAC9) on the Knowledge Base page at Dell.com

Systems Management software support matrix

Table 32. Systems Management software support matrix

Categories	Features	PE mainstream
Embedded Management and In-band	iDRAC9 (Express, Enterprise, and Datacenter licenses)	Supported
Services	OpenManage Mobile	Supported
	OM Server Administrator (OMSA)	Supported
	iDRAC Service Module (iSM)	Supported
	Driver Pack	Supported
Change Management	Update Tools (Repository Manager, DSU, Catalogs)	Supported
	Server Update Utility	Supported
	Lifecycle Controller Driver Pack	Supported
	Bootable ISO	Supported
Console and Plug-ins	OpenManage Enterprise	Supported
	Power Manager Plug-in	Supported
	Update Manager Plug-in	Supported
	SupportAssist Plug-in	Supported
	CloudIQ	Supported
Integrations and connections	OM Integration with VMware Vcenter/vROps	Supported
	OM Integration with Microsoft System Center (OMIMSC)	Supported
	Integrations with Microsoft System Center and Windows Admin Center (WAC)	Supported



Table 32. Systems Management software support matrix (continued)

Categories	Features	PE mainstream
	ServiceNow	Supported
	Ansible	Supported
	Supported	
Security	Secure Enterprise Key Management	Supported
	Secure Component Verification	Supported
Standard operating system	Red Hat Enterprise Linux, SUSE, Windows Server 2019 or 2022, Ubuntu, CentOS	Supported (Tier-1)

54 Dell OpenManage Systems Management

Inserido ao protocolo **22.951.206-4** por: **Pedro Henrique Golin Linhares** em: 31/03/2025 12:49. A autenticidade deste documento pode ser validada no endereço: https://www.eprotocolo.pr.gov.br/spiweb/validarDocumento com o código: **692a5ef82beb173ed33dad4ee6505546**.



Appendix A. Standards compliance

The system conforms to the following industry standards.

Table 33. Industry standard documents

Standard	URL for information and specifications	
ACPI Advance Configuration and Power Interface Specification, v6.4	Uefi specifications and tools	
Ethernet IEEE Std 802.3-2022	ieee standards	
MSFT WHQL Microsoft Windows Hardware Quality Labs	microsoft.com/whdc/system/platform/pcdesign/desguide/ serverdg.mspx	
IPMI Intelligent Platform Management Interface, v2.0	intel.com/design/servers/ipmi	
DDR5 Memory DDR5 SDRAM Specification	jedec.org/standards-documents/docs/jesd79-4.pdf	
PCI Express PCI Express Base Specification, v5.0	pcisig.com/specifications/pciexpress	
PMBus Power System Management Protocol Specification, v1.2	pmbus specification and revisions	
SAS Serial Attached SCSI, 3 (SAS-3) (T10/INCITS 519)	SCSI storage interfaces information	
SATA Serial ATA Rev. 3.3	sata-io.org page	
SMBIOS System Management BIOS Reference Specification, v3.3.0	BIOS reference specification page	
TPM Trusted Platform Module Specification, v1.2 and v2.0	trustedcomputinggroup org page	
UEFI Unified Extensible Firmware Interface Specification, v2.7	UEFIF specifications	
PI Platform Initialization Specification, v1.7		
USB Universal Serial Bus v2.0 and SuperSpeed v3.0 (USB 3.1 Gen1)	USB Implementers Forum, Inc. USB document library	
NVMe Express Base Specification. Revision 2.0c	NVME specifications	
 NVMe Command Set Specifications NVM Express NVM Command Set Specification. Revision 1.1c NVM Express Zoned Namespaces Command Set. Revision 1.0c NVM Express® Key Value Command Set. Revision 1.0c 		
 NVMe Transport Specifications 1. NVM Express over PCle Transport. Revision 1.0c 2. NVM Express RDMA Transport Revision. 1.0b 3. NVM Express TCP Transport. Revision 1.0c 		
NVMe NVM Express Management Interface. Revision 1.2c		
NVMe NVMe Boot Specification. Revision 1.0		



Appendix B: Additional specifications

Topics:

- Chassis dimensions
- System weight
- Video specifications
- USB ports specifications
- Environmental specifications

Chassis dimensions



Figure 41. Chassis dimensions

Table 34. PowerEdge R760xs chassis dimensions

Xa	ХЬ	Y	Za	Zb	Zc
482.0 mm (18.97 inches)	434.0 mm (17.08 inches)	86.8 mm (3.41 inches)	22.0 mm (0.86 inches) Without bezel	677.44 mm (26.67 inches) Ear to L bracket housing	685.78 mm (26.99 inches) Ear to PSU handle without velcro strap

56 Appendix B: Additional specifications

Inserido ao protocolo **22.951.206-4** por: **Pedro Henrique Golin Linhares** em: 31/03/2025 12:49. A autenticidade deste documento pode ser validada no endereço: https://www.eprotocolo.pr.gov.br/spiweb/validarDocumento com o código: 692a5ef82beb173ed33dad4ee6505546.



Table 34. PowerEdge R760xs chassis dimensions

Xa	Xb	Y	Za	Zb	Zc
			35.84 mm (1.41 inches) With bezel	650.24 mm (25.6 inches) Ear to PSU surface	

(i) NOTE: Zb is the nominal rear wall external surface where the system board I/O connectors reside.

System weight

Table 35. PowerEdge R760xs systemweight

System configuration	Maximum weight (with all drives/SSDs/bezel)
16+8 x 2.5-inch	25.92 kg (57.14 lb)
16 x 2.5-inch	24.58 kg (54.18 lb)
12 x 3.5-inch	28.82 kg (63.53 lb)
8 x 3.5-inch	25.84 kg (54.96 lb)
8 x 2.5-inch	21.56 kg (47.53 lb)
No backplane configuration	19.40 kg (42.76 lb)

Video specifications

The PowerEdge R760xs system supports integrated Matrox G200 graphics controller with 16 MB of video frame buffer.

Table 36. Supported video resolution options for the system

Resolution	Refresh rate (Hz)	Color depth (bits)
1024 x 768	60	8, 16, 32
1280 x 800	60	8, 16, 32
1280 x 1024	60	8, 16, 32
1360 x 768	60	8, 16, 32
1440 x 900	60	8, 16, 32
1600 x 900	60	8, 16, 32
1600 x 1200	60	8, 16, 32
1680 x 1050	60	8, 16, 32
1920 x 1080	60	8, 16, 32
1920 x 1200	60	8, 16, 32



USB ports specifications

Table 37. PowerEdge R760xs USB specifications

Front		Rear		Internal (Optional)	
USB port type	No. of ports	USB port type	No. of ports	USB port type	No. of ports
USB 2.0- compliant port	One	USB 2.0- compliant port	One	Internal USB 3.0- compliant port	One
iDRAC Direct port (Micro-AB USB 2.0- compliant port)	One	USB 3.0- compliant port	One		

(i) NOTE: The micro USB 2.0 compliant port can only be used as an iDRAC Direct or a management port.

NOTE: The USB 2.0 specifications provide a 5 V supply on a single wire to power connected USB devices. A unit load is defined as 100 mA in USB 2.0, and 150 mA in USB 3.0. A device may draw a maximum of 5 unit loads (500 mA) from a port in USB 2.0; 6 (900 mA) in USB 3.0.

NOTE: The USB 2.0 interface can provide power to low-power peripherals but must adhere to USB specification. An external power source is required for higher-power peripherals to function, such as external CD/DVD Drives.

Environmental specifications

NOTE: For additional information about environmental certifications, see the Product Environmental Datasheet that are located with the Documentation on https://www.dell.com/support.

Table 38. Continuous Operation Specifications for ASHRAE A2

Temperature, humidity and, operational altitude	Allowable continuous operations
Temperature range for altitudes <= 900 m (<= 2953 ft)	10-35°C (50-95°F) with no direct sunlight on the equipment
Humidity percent range (non-condensing at all times)	8% RH with -12°C (10.4°F) minimum dew point to 80% RH with 21°C (69.8°F) maximum dew point
Operational altitude derating	Maximum temperature is reduced by 1°C/300 m (1.8°F/984 Ft) above 900 m (2953 Ft)

Table 39. Continuous Operation Specifications for ASHRAE A3

Temperature, humidity and, operational altitude	Allowable continuous operations
Temperature range for altitudes <= 900 m (<= 2953 ft)	5–40°C (41–104°F) with no direct sunlight on the equipment
Humidity percent range (non-condensing at all times)	8% RH with -12°C (10.4°F) minimum dew point to 85% RH with 24°C (75.2°F) maximum dew point
Operational altitude derating	Maximum temperature is reduced by 1°C/175 m (1.8°F/574 Ft) above 900 m (2953 Ft)

Table 40. Continuous Operation Specifications for ASHRAE A4

Temperature, humidity and, operational altitude	Allowable continuous operations
Temperature range for altitudes <= 900 m (<= 2953 ft)	5-45°C (41-113°F) with no direct sunlight on the equipment

58 Appendix B: Additional specifications


Table 40. Continuous Operation Specifications for ASHRAE A4 (continued)

Temperature, humidity and, operational altitude	Allowable continuous operations
Humidity percent range (non-condensing at all times)	8% RH with -12°C (10.4°F) minimum dew point to 90% RH with 24°C (75.2°F) maximum dew point
Operational altitude derating	Maximum temperature is reduced by 1°C/125 m (1.8°F/410 Ft) above 900 m (2953 Ft)

Table 41. Continuous Operation Specifications for Rugged Environment

Temperature, humidity and, operational altitude	Allowable continuous operations
Temperature range for altitudes <= 900 m (<= 2953 ft)	5-45°C (41-113°F) with no direct sunlight on the equipment
Humidity percent range (non-condensing at all times)	8% RH with -12°C (10.4°F) minimum dew point to 90% RH with 24°C (75.2°F) maximum dew point
Operational altitude derating	Maximum temperature is reduced by 1°C/125 m (1.8°F/410 Ft) above 900 m (2953 Ft)

Table 42. Common Environmental Specifications for ASHRAE A2, A3, A4, and Rugged

Temperature, humidity and, operational altitude	Allowable continuous operations
Maximum temperature gradient (applies to both operation and non-operation)	20°C in an hour* (36°F in an hour) and 5°C in 15 minutes (9°F in 15 minutes), 5°C in an hour* (9°F in an hour) for tape (i) NOTE: * - Per ASHRAE thermal guidelines for tape hardware, these are not instantaneous rates of temperature change.
Non-operational temperature limits	-40°C to 65°C (-104°F to 149°F)
Non-operational humidity limits	5% to 95% RH with 27°C (80.6°F) maximum dew point
Maximum non-operational altitude	12,000 meters (39,370 feet)
Maximum operational altitude	3,048 meters (10,000 feet)

Table 43. Maximum vibration specifications

Maximum vibration	Specifications		
Operating	0.21 G_{rms} at 5 Hz to 500 Hz for 10 minutes (all operation orientations)		
Storage	1.88 G_{rms} at 10 Hz to 500 Hz for 15 minutes (all six sides tested)		

Table 44. Maximum shock pulse specifications

Maximum shock pulse	Specifications
Operating	Six consecutively performed shock pulses in the positive and negative x, y, and z axis of 6 G for up to 11 ms.
Storage	Six consecutively performed shock pulses in the positive and negative x, y, and z axis (one pulse on each side of the system) of 71 G for up to 2 ms.

The following table defines the limitations that help avoid any equipment damage or failure from particulates and gaseous contamination. If the levels of particulate or gaseous pollution exceed the specified limitations and result in equipment damage or failure, you may need to rectify the environmental conditions. Re-mediation of environmental conditions is the responsibility of the customer.



Table 45. Particulate and gaseous contamination specifications

Particulate contamination	Specifications			
Air filtration	 Data center air filtration as defined by ISO Class 8 per ISO 14644-1 with a 95% upper confidence limit. (i) NOTE: The ISO Class 8 condition applies to data center environments only. This air filtration requirement does not apply to IT equipment designed to be used outside a data center, in environments such as an office or factory floor. (i) NOTE: Air entering the data center must have MERV11 or MERV13 filtration. 			
Conductive dust	 Air must be free of conductive dust zinc whiskers, or other conductive particles. NOTE: This condition applies to data center and non-data center environments. 			
Corrosive dust	Air must be free of corrosive dust. Residual dust present in the air must have a deliquescent point less than 60% relative humidity. (i) NOTE: This condition applies to data center and non-data center environments.			

Table 46. Gaseous contamination specifications

Gaseous contamination	Specification		
Copper coupon corrosion rate	<300A/month per class G1 as defines by ANSI/ISA71.04-2013		
Silver coupon corrosion rate	<200A/month as defined by ANSI/ISA71.04-2013		

Thermal restriction matrix

Table 47. Processor and heat sink matrix

Heat sink	Processor TDP	
STD HSK	< 185 W CPU SKUs	
HPR HSK	185 W-250 W CPU SKUs (12 x 3.5-inch drive configuration not supported)	
	125W-250W CPU SKUs (12 x 3.5-inch drive configuration supported)	

Table 48. Label reference

Label	Description		
STD	Standard		
HPR (Silver)	High Performance Silver (HPR) fan		
HPR (Gold)	High Performance Gold (VHP) fan		
HSK	Heat sink		

() NOTE: The ambient temperature of the configuration is determined by the critical component in that configuration. For example, if the processor's supported ambient temperature is 35°C, the DIMM is 35°C, and the GPU is 30°C, the combined configuration can only support 30°C.

60 Appendix B: Additional specifications



Table 49. Supported ambient temperature for processors for R760xs

R760xs	R760xs									
configur	ation		No backpla ne	8 x 3.5- inch SAS configu ration	12 x 3.5- inch SAS configu ration	12 x 3.5- inch configu ration with rear drive module	8 x 2.5- inch SAS configu ration	8 x 2.5- inch NVMe configu ration	16 x 2.5- inch SAS configu ration	16 x 2.5- inch + 8 x 2.5- inch NVMe configu ration
EMR	4514Y	150 W	45°C	40°C	35°C	35°C	40°C	40°C	40°C	40°C
CPU	5512U	185 W	45°C	40°C	35°C	35°C	40°C	40°C	40°C	40°C
	6526Y/6534	195 W	35°C	35°C	35°C	35°C	35°C	35°C	35°C	35°C
	6548Y+/ 6542Y/6548N	250 W	35°C	35°C	30°C	30°C	35°C	35°C	35°C	35°C
SPR	4509Y	125 W	45°C	40°C	35°C	35°C	40°C	40°C	40°C	40°C
CPU	4510	150 W	45°C	40°C	35°C	35°C	40°C	40°C	40°C	40°C
SPR	3408U	125 W	45°C	40°C	35°C	35°C	40°C	40°C	40°C	40°C
MCC CPU	5416S/ 4410T/ 4410Y/5415+	150 W	45°C	40°C	35°C	35°C	40°C	40°C	40°C	40°C
	4416	165 W	45°C	40°C	35°C	35°C	40°C	40°C	40°C	40°C
	5418Y/ 5412U/6426Y	185 W	45°C	40°C	35°C	35°C	40°C	40°C	40°C	40°C
	5420+/ 6438Y+	205 W	35°C	35°C	35°C	35°C	35°C	35°C	35°C	35°C
	6448Y/6442Y	225 W	35°C	35°C	35°C	35°C	35°C	35°C	35°C	35°C
	6414U	250 W	35°C	35°C	30°C	30°C	35°C	35°C	35°C	35°C
Memory	96 GB RDIMM 5200	8.1 W, 1DPC	35°C	35°C	30°C	30°C	35°C	35°C	35°C	35°C
	64 GB RDIMM 5200	7.7 W, 1DPC	45°C	40°C	35°C	35°C	40°C	40°C	40°C	40°C
	32 GB RDIMM 5200	5.1 W, 1DPC	45°C	40°C	35°C	35°C	40°C	40°C	40°C	40°C
	64 GB RDIMM 4800	12 W, 1DPC	45°C	40°C	35°C	35°C	40°C	40°C	40°C	40°C
	32 GB RDIMM 4800	10 W, 1DPC	45°C	40°C	35°C	35°C	40°C	40°C	40°C	40°C
PCle		45°C	40°C	35°C1	35°C1	40°C	40°C	40°C	40°C	
A2 GPU ⁶			35°C	30°C	Not supporte d	Not supporte d	35°C	35°C	30°C	30°C
OCP			45°C	40°C	35°C ²	35°C ²	40°C	40°C	40°C	40°C
BOSS			35°C	35°C	35°C	35°C	35°C	35°C	35°C	35°C

(i) NOTE:

- 1. Max supported thermal tier of PCIe card is Tier 5.
- **2.** Max supported thermal tier of OCP is Tier 5.



- HPR Sliver fan is required from fan zone 2 to fan zone 6 for 8 x 2.5-inch NVMe, 16 x 2.5-inch SAS/SATA + 8 x 2.5-inch NVMe, 12 x 3.5-inch drives or GPU configurations.
- 4. Optional fan zone 1 has to be populated with HPR Gold fan is for BOSS, GPU or rear drive module populations.
- 5. PCIe slot priority of Nvidia A2 GPU is constrained on slot #3, #4, #6.
- 6. HPR heatsink is required for ≥ 185 W CPUs, 12 x 3.5-inch drives or 12 x 3.5-inch drives with rear storage module configurations.
- 7. DIMM blank is required for 12 x 3.5-inch SAS/SATA with rear storage module.
- 8. Fan blank is required on fan zone 1 when no fan population.
- **9.** OCP shroud is required for OCP card population without PCIe riser module installed.
- 10. CPU blank is required for single processor configuration.
- 11. Rear drive module does not support Kioxia CM6 series, Samsung PM1735 series, Hynix PE8010/PE8110 ≥ 7.68 TB,
- Samsung PM1733a > 1.92 TB, Samsung PM1735a > 1.6 TB and Redtail NVMe drive.

(i) NOTE: The fan speed in the 3.5-inch chassis is limited to 90% due to the drive dynamic profile.

Table 50. Fan population rule for R760xs

configuratio n	No backplane	8 x 3.5-inch SAS	12 x 3.5- inch SAS	8 x 2.5-inch SAS	8 x 2.5-inch NVMe	16 x 2.5-inch SAS	24 x 2.5-inch (16 x 2.5-inch
Optional HW							NVMe)
Default	Fan 2 to Fan 6 with STD fan	Fan 2 to Fan 6 with STD fan	Fan 2 to Fan 6 with HPR Silver fan	Fan 2 to Fan 6 with STD fan	Fan 2 to Fan 6 with HPR Silver fan	Fan 2 to Fan 6 with STD fan	Fan 2 to Fan 6 with HPR Silver fan
Rear Module	Not supported	Not supported	Fan 1 with HPR Gold fan	Not supported	Not supported	Not supported	Not supported
			Fan 2 to Fan 6 with HPR Silver fan				
BOSS N1	Fan 1 with HPR Gold fan	Fan 1 with HPR Gold fan	Fan 1 with HPR Gold fan				
	Fan 2 to Fan 6 with STD fan	Fan 2 to Fan 6 with STD fan	Fan 2 to Fan 6 with HPR Silver fan	Fan 2 to Fan 6 with STD fan	Fan 2 to Fan 6 with HPR Silver fan	Fan 2 to Fan 6 with STD fan	Fan 2 to Fan 6 with HPR Silver fan
GPU	Fan 1 with HPR Gold fan	Fan 1 with HPR Gold fan	Not supported	Fan 1 with HPR Gold fan	Fan 1 with HPR Gold fan	Fan 1 with HPR Gold fan	Fan 1 with HPR Gold fan
	Fan 2 to Fan 6 with HPR Silver fan	Fan 2 to Fan 6 with HPR Silver fan		Fan 2 to Fan 6 with HPR Silver fan	Fan 2 to Fan 6 with HPR Silver fan	Fan 2 to Fan 6 with HPR Silver fan	Fan 2 to Fan 6 with HPR Silver fan

Thermal Restrictions for PCIe adapter NIC and other network cards with iDRAC

- Cannot support PCIe cards with the cooling requirement more than 300LFM at 55C in a 12 x 3.5-inch SAS/SATA configuration.
- Solarflare Melrose DP 25 GBE SFP28 (TTKWY) not supported with 12 x 3.5-inch SAS/SATA configuration.
- 100 Gb network adapter or 100 Gb OCP is not supported in the 12 x 3.5-inch SAS/SATA configuration.
- Few 25 GB OCP cards with the cooling requirement more than 250LFM at 55C (3Y64D/4TRD3 / GGGDF/R1KTR / Y4VV5) is not supported in 12x3.5" SAS/SATA configuration.
- The 12 x 3.5-inch SAS/SATA configuration requires the optical transceiver with higher temperature spec (\geq 85°C) to support (M14MK / N8TDR).



- Quad port OCP (3Y64D/Y4VV5) requires the optical transceiver with higher temperature spec (≥ 85°C) to support (M14MK).
- 100 Gb network adapter cannot support the transceivers as 14NV5/9JKK2 / QSFP56 (MFS1S00-VxxxE/HxxxE).
- The H965e is limited to populate in PCI slot 3 in a 12 x 3.5-inch SAS/SATA configuration.
- Mellanox CX7 NDR200 card has few limitations of PCI slot locations.

Table 51. Mellanox CX7 NDR200 slot location limitations

Storage configuration	Slots on 3.5-inch configuration	Slots on 2.5-inch configuration
Gen5 PCIe sloit support for CX7 NDR200	3, 4	3, 4
Gen4 PCIe sloit support for CX7 NDR200	6	5, 6

Thermal restrictions for extended ambient support (ASHRAE A3/A4)

- Two PSUs are required in redundant mode. Single PSU failure is not supported.
- 12 x 3.5-inch SAS/SATA configuration is not supported.
- BOSS(M.2) module is not supported.
- CPU TDP > 185 W is not supported.
- PCle card TDP > 25 W is not supported.
- OCP cards with transmission rate higher than 25 GB is not supported.
- OCP transceiver spec $\leq 75^{\circ}$ C is not supported.
- 8 x 3.5-inch SAS/SATA, 8 x 2.5-inch SAS/SATA, 8 x 2.5-inch NVMe, 16 x 2.5-inch SAS/SATA, 16 x 2.5-inch SAS/SATA + 8x 2.5-inch NVMe configurations are limited to support A3.
- 128 GB+ memory is not supported .
- The rear drive is not supported.



Appendix C Additional resources

Table 52. Additional resources

Resource	Description of contents	Location
Installation and Service Manual	 This manual, available in PDF format, provides the following information: Chassis features System Setup program System indicator codes System BIOS Remove and replace procedures 	Dell.com/Support/Manuals
	DiagnosticsJumpers and connectors	
Getting Started Guide	This guide ships with the system, and is also available in PDF format. This guide provides the following information: Initial setup steps	Dell.com/Support/Manuals
Rack Installation Guide	This document ships with the rack kits, and provides instructions for installing a server in a rack.	Dell.com/Support/Manuals
System Information Label	The system information label documents the system board layout and system jumper settings. Text is minimized due to space limitations and translation considerations. The label size is standardized across platforms.	Inside the system chassis cover
QR code for system resources	This code on the chassis can be scanned by a phone application to access additional information and resources for the server, including videos, reference materials, service tag information, and Dell contact information.	Inside the system chassis cover
Enterprise Infrastructure Planning Tool (EIPT)	The Dell online EIPT enables easier and more meaningful estimates to help you determine the most efficient configuration possible. Use EIPT to calculate the power consumption of your hardware, power infrastructure, and storage.	Dell.com/calc

Topics:

- Documentation
- Customer kits



Documentation

This section provides information about the documentation resources for your system.

Table 53. Documentation resources

Document	Location
Factory Configuration Matrix	Sales Portal
SPM (Slot Priority Matrix)	Sales Portal
NDA Deck	Sales Portal
Installation and Service Manual (ISM)	https://www.dell.com/poweredgemanuals
Field Service Manual (FSM)	https://www.dell.com/poweredgemanuals > Sing in
Technical Guide	Dell.com > Product page > Product Details
Spec Sheet	Dell.com > Product page > Product Details

Customer kits

Dell Upgrades

It is not always possible to plan for new applications, future workloads, and business needs. Unleash the full power of your Dell Technologies Infrastructure. When budget does not permit the purchase of new servers, Dell Upgrades is a cost-effective method to repurpose and unleash the full power of existing server, storage, and networking infrastructure.

- Protect your mission-critical operations by using only genuine Dell OEM-validated Upgrades and the technical expertise of Dell ProSupport
- Flex and scale existing infrastructure by upgrading, adding memory or storage drives to cost-effectively and quickly meet new workloads and demands
- Dell Upgrades are the same peripheral commodities that your customer may improve or maintain their server after the initial point of sale

Upgrades portfolio

Table 54. Upgrade category

Dell Upgrade Category	Sample Picture	Dell Upgrade Category Offerings
Memory Memory upgrades are essential for keeping your customers operating at peak performance as their business needs grow and their workloads increase. We tend to see strong demand for server memory because it is the easiest and most cost-effective way to improve system performance.	Contraction of the second seco	DDR5 5600 MT/s and 4800 MT/s



Table 54. Upgrade category (continued)

Dell Upgrade Category	Sample Picture	Dell Upgrade Category Offerings
Storage Dell offers both solid-state drive and hard disk drive storage options for enterprise systems with SATA, SAS or NVMe interfaces. SSDs excel in speed, high-performance I/O requirements, and high reliability due to the lack of spinning disks. Hard Disk Drives (HDDs) store data on spinning disks and offer value for the amount of data storage for the price. Dell offers both solid-state drive and hard disk drive storage options for enterprise systems with SATA, SAS interfaces. SSDs excel in speed, high- performance I/O requirements, and high reliability due to the lack of spinning disks. Hard Disk Drives (HDDs) store data on spinning disks and offer value for the amount of data storage for the price.		HDD: SATA, SAS interface SSD: SATA, SAS, PCI NVMe interface Tape Drive or Media
Processor Processor upgrades help customers perform and accomplish more tasks overall, saving them valuable time. Our processor upgrades include Intel® Xeon® Scalable processors to meet your customers workload needs with increased cores and improved security.		Processors (Intel) Heat sinks
Networking and Optics Our networking and optics components —network interface cards, transceivers, optical cables, and more—are key in today's data center environment, helping customers to improve bandwidth to better manage increase in workloads, devices, users, and interconnected systems.		Network cards Transceivers (Optics)



Table 54. Upgrade category (continued)

Dell Upgrade Category	Sample Picture	Dell Upgrade Category Offerings
Accessories: Dell sells accessories like	A	Controller cards
bezels, controller cards, GPU, PERC and		Power supplies
other components to complete the Dell Upgrades portfolio and redundancies.		Cables
	· .	Rail kits
	E Hannah	Bezels
		Power cords
		GPU
	Dellever	PERC
	80	BOSS
		Power cords
	States	Cable Management Arm (CAM)
	Sector -	Fans
	10001 34	Serial board
	NOT OCC	Internal USB
	2800W	

Upgrades reference links

- Main Upgrades Page
- Customer Kit Selector
- Dell Parts Finder Tool (Customer Facing Tool)



Appendix D: Service and support

Topics:

- Why attach service contracts
- ProSupport Infrastructure Suite
- Specialty Support Services
- Dell deployment services
- Supplemental Deployment Services
- Unique Deployment Scenarios
- DAY 2 Automation Services with Ansible
- Dell Technologies Consulting Services

Why attach service contracts

Dell PowerEdge servers include a standard hardware warranty that highlights our commitment to product quality by guaranteeing repair or replacement of defective components. While industry-leading, our warranties are limited to 1 or 3 years, depending on model, and do not cover software assistance. Call records show that failure rates for servers are roughly 1% and more commonly, customers seek Dell technical support for software-related issues like configuration guidance, troubleshooting, upgrade assistance, or performance tuning. Encourage customers to purchase ProSupport service contracts to supplement warranty coverage and ensure optimal support for both hardware and software. ProSupport provides a complete hardware guarantee beyond the original warranty period (up to 12 years: including seven years standard support and an additional five years of Post-Standard Support). Details of the ProSupport Suite and benefits are listed below.

ProSupport Infrastructure Suite

ProSupport Infrastructure Suite is a set of support services that enable customers to build the solution that is right for their organization. It is an industry-leading, enterprise-class support that aligns with the criticality of your systems, the complexity of your environment, and the allocation of your IT resources.



ProSupport Infrastructure Suite | Enhanced value across all offers!

	Basic Herdware Support	ProSupport For Intrastructure	ProSupport Plus for Infrastructure	Changes with August 2025 minute
Technical support availability and response objective	Why intermediated	247, mmedate	247, mmodiate	No shange
Covered producte	Hintkeara	Hastwere & Sofware	Hartheare & Boffware	No shanga
Onsite response service level	1000	NSO or 4-hour	4-hour	Projugatori Plus NBD is refined.
ProSupport AlOps platforms			•	McGenuce300 and TechDirect (all offen) Close90 (PerSupport & PerSupport Plus)
Del: Security Advisories	(11)		•	Available on additional products
Proactive issue detection with automated case creation		•	•	New to Desc
Predictive hardware anomaly detection		6 (•	New In ProSupport
Access to software updates		1	•	No change
GoudiQ health and cybersecurity monitoring & analytics			•	Enternant Instalium
Incident Manager for Severity 1 cases		1 (Dec	•	No chenge
Mission Official support			•	Eliforcos fontem
Priority access to remote senior support engineers1				No change
Service Account Manager			•	No change
Proactive system maintenance				No change
Limited 3 rd party software support ^{er}				No coverge

*Software looses can be pumbased through Dell or BYOL - see Service Descriptions for details

DELL'actuations

Figure 42. ProSupport Enterprise Suite

ProSupport Plus for Infrastructure

ProSupport Plus for Infrastructure is the ultimate solution for customers seeking preventative maintenance and optimal performance on their business-critical assets. The service caters to customers who require proactive, predictive, and personalized support for systems that manage critical business applications and workloads. When customers purchase PowerEdge server, we recommend ProSupport Plus, our proactive and preventative support service for business-critical systems. ProSupport Plus provides all the benefits of ProSupport, including the following "Top five reasons to buy ProSupport Plus (PSP)"

- 1. **Priority access to specialized support experts:** Immediate, advanced troubleshooting from an engineer that understands Dell infrastructure solutions.
- 2. Mission Critical Support: When critical (Severity 1) support issues happen, the customer is assured that we do all that we can to get them back up and running as quickly as possible.
- 3. Service Account Manager: A customer's #1 support advocate, ensuring they get the best possible proactive and predictive support experience.
- 4. Systems maintenance: On a semiannual basis, we will keep a customer's ProSupport Plus system(s) up to date by installing the latest firmware, BIOS, and driver updates to improve performance and availability.
- 5. Third-party software support: Dell is a customer's single point of accountability for any eligible third-partysoftware that is installed on their ProSupport Plus system, whether they purchased the software from us or not.

ProSupport for Infrastructure

Comprehensive 24x7 support for hardware and software – best for production, but not critical, workloads and applications. The ProSupport service offers highly trained experts around the clock and around the globe to address IT needs. We help minimize disruptions and maximize availability of PowerEdge server workloads with:

- 24x7 support through phone, chat and online
- A central point of accountability for all hardware and software issues
- Hypervisor, operating system and application support
- Dell security advisories
- Onsite response service levels 4 hour or Next Business Day options
- Proactive issue detection with automated case creation



- Predictive hardware anomaly detection
- Incident Manager assigned for Severity 1 cases
- Collaborative third-party support
- Access to AIOps Platforms (MyService360, TechDirect, and CloudIQ)
- Consistent experience regardless of where customers are located or what language that they speak.

Basic Hardware Support

Provides reactive hardware support during normal business hours, excluding local national holidays. No software support orsoftware-related guidance. For improved levels of support, choose ProSupport or ProSupport Plus.

Specialty Support Services

Optional specialty support services complement the ProSupport Infrastructure Suite to provide additional proficiencies that are critical for modern data center operations.

Hardware coverage add-ons to ProSupport

• Keep Your Hard Drive (KYHD), Keep Your Component (KYC), or Keep Your GPU:

Normally if a device fails under warranty, Dell replaces it using a one-for-one exchange process.KYHD/KYCC/KYGPU gives you the option to retain your device. It provides full control of sensitive data and minimizes security risk by letting you retain possession of failed drives, components, or GPU when receiving replacement parts without incurring additional cost.

• Onsite Diagnosis Service:

Ideal for sites with non-technical staff. Dell field technician performs initial troubleshooting diagnosis onsite and transfers to Dell remote engineers to resolve the issue.

• ProSupport Add-on for HPC:

Sold as an add-on to a ProSupport service contract, the ProSupport Add-on for HPC provides solution-aware support to cover the additional requirements that are required to maintain an HPC environment such as:

- \circ $\,$ Access to senior HPC experts
- Advanced HPC cluster assistance: performance, interoperability, and configuration
- Enhanced HPC solution level end-to-end support
- Remote pre-support engagement with HPC Specialists during ProDeploy implementation
- ProSupport Add-on for Telco (Respond & Restore):

An add-on service designed for the top 31 TELCO customers globally, Respond & Restore provides direct access to Dell solution experts who specialize in TELCO carrier-grade support. This add-on also provides a hardware uptime guarantee, meaning if a system fails, Dell has it installed and operational within 4 hours for Severity 1 issues. Dell incurs penalties and fees if SLAs are not met.

Personalized Support and Supplemental Site-wide Expertise

• Technical Account Manager:

Designated technology lead who monitors and manages the performance and configuration of specific technology sets.

• Designated Remote Support:

Personalized support expert who manages all troubleshooting and resolution of IT assets.

• Multivendor Support Service:

Support your third-party devices as one service plan for servers, storage, and networking (includes coverage for: Broadcom, Cisco, Fujitsu, HPE, Hitachi, Huawei, IBM, Lenovo, NetApp, Oracle, Quanta, SuperMicro and others).



Services for large enterprises

• ProSupport One for Data Center:

ProSupport One for Data Center offers flexible site-wide support for large and distributed data centers with more than 1,000 assets (combined total of server, storage, networking, so forth). This offering is built on standard ProSupport features that leverage our global scale and are tailored to specific customer needs. While not for everyone, this service option offers a truly unique solution for our largest customers with the most complex environments.

- Team of assigned Services Account Managers with remote or onsite options
- Assigned technical and field engineers who are trained on the customer's environment and configurations.
- On-demand reporting and recommendations that are enabled by ProSupport AlOps tools (MyService360, TechDirect, and CloudIQ)
- Flexible onsite support and parts options that fit their operational model
- A tailored support plan and training for their operations staff
- ProSupport One for CSPs (Cloud Serviced Providers)

ProSupport One for CSPs is a unique offer that is designed for a limited set of Dell accounts purchasing Gen Al computing solutions greater than 1,000 servers and \$250M in sales. PS1 for CSPs improves the entire services experience combining support, deployment (rack integration), residency services, a designated support engineer and the LOIS parts locker as one holistic bundle. Special pricing has been determined to compete effectively against competitors and provide the best customer experience. PS1 for CSPs can only be sold with XE Servers and all networking platforms (Dell and NVIDIA). All other products would be eligible for the standard PS1DC not this special bundle offer. More details on PS1 for CSPs here.

• Logistics Online Inventory Solution (LOIS)

Ideal for large organizations that have their own staff to support their data center. Dell offers a service that is called Logistics Online Inventory Solution which is an onsite parts locker that provides self-maintainers with a local inventory of common replacement components. Having access to these parts lockers allows the self-maintainer to replace a failed component immediately without delay. Each replacement part would automatically initiate a replenishment of the parts inventory that is shipped the next day or delivered onsite by Dell during a regular scheduled visit (called Scheduled Onsite Service). As part of the LOIS system, customers can integrate their systems directly to Dell TechDirect using APIs to help streamline the support management process.

End-of-Life Services

• Post Standard Support (PSS)

Extend service life beyond the initial seven years of ProSupport, adding up to five more additional years of hardware coverage.

Data Sanitization & Data Destruction

Renders data unrecoverable on repurposed or retired products, ensuring security of sensitive data and enabling compliance and provides NIST-compliant certification.

• Asset Recovery Services

Recycle, resale, and disposal of hardware. Helps you securely and responsibly retire IT assets that are no longer needed while protecting both your business and the planet.

Dell deployment services

Dell ProDeploy Infrastructure Suite

ProDeploy Infrastructure Suite provides a variety of deployment offerings that satisfy a customer's unique needs. It is made up of 5 offers: ProDeploy Configuration Services, ProDeploy Rack Integration Services, Basic Deployment, ProDeploy, and ProDeploy Plus.

ProDeploy Infrastructure Suite for servers

Versatile choices for accelerated deployments



FIS. <u>932</u> Moy. 32

Figure 43. ProDeploy Infrastructure Suite for servers

The new Factory Services consist of two tiers of deployment that happen prior to shipping to the customer's site.

Factory Based Services:

- ProDeploy Factory Configuration Ideal for customers buying servers in volume and seeking pre-configuration prior to shipping such as: custom image, system settings, and asset tagging so it arrives ready to use out of the box. Furthermore, servers can be packaged and bundled to meet specific shipping and distribution requirements for each customer location to facilitate the rollout process. Upsell one of the field based services (below) if a customer needs assistance with the final server installation.
- ProDeploy Rack Integration Ideal for customers seeking to build out fully integrated racks prior to shipping. These rack builds include hardware install, cabling, and full system configuration. You can also add-on a factory stress test and optional on-site final rack configuration to complete the rack installation.
 - STANDARD SKUs for Rack Integration is available in US only and requires:
 - 20 or more devices (R and C series servers and all Dell or non-Dell switches). Use Informational SKUs for Dell switches or 3rd party products
 - Shipping to contiguous US
 - USE CUSTOM QUOTE for Rack Integration for:
 - All countries except USA
 - Racks containing less than 20 servers
 - Any rack that includes VxRail or Storage
 - Shipping outside contiguous US
 - Shipping to multiple locations

Field Based Services:

- Basic Deployment consists of the hardware installation, cabling and firmware update during normal standard business hours. Basic Deployment is traditionally sold to Competency Enabled Partners. Competency enabled partners often have Dell do the hardware installation while they complete the software configuration.
- ProDeploy consists of your hardware installation and configuration of the software using offshore resources. ProDeploy is great for customers who are price sensitive or who are remote from their data centers and don't require an onsite presence.
- ProDeploy Plus will give you in-region or onsite resources to complete the engagement for the customer. It also comes with additional features such as Post Deployment Configuration Assistance and Training Credits.



ProDeploy Infrastructure Suite | Factory services

		Pactori asaco scrute s	
		ProDeployFactory Configuration	ProDeploy Rack Integration
and the second sec	Single point of contact for project management	0/	•
	RAID, BIOS and IDRAC configuration		
sset configuration	Firmware freeze		
	Asset Tagging and Reporting		
	Customer system image		0
	Site readiness review and implementation planning		
	Hardware racking and cabling		
actory implementation	SAM engagement for ProSupport Plus entitled accounts/devices	2	
	Deployment verification, documentation, and knowledge transfer		(
	White glove logistics		
	Onsite final configuration		Onsite add-on
elivery	Install support software and connect with Dell Technologies		Onsite add-on
	Basic Deployment	Optional onsite installation	
	Online collaborative environment for elements, managing and tracking delivery		•

Figure 44. ProDeploy Infrastructure Suite - Factory services

ProDeploy Infrastructure Suite | Field services

		Basic Deployment	ProDeptoy	ProCepio Plas
	Single point of contact for project management.	•		In-region
	Etto readmoso review			•
e.o.coleniane	Implementation planning*	- II	•	
	SAM engagement for ProSupport Plus entitled devices	2	- 10 10	•
	Deployment service hours	Business hours	24.07	24x7
Deploym Oncils h herower	Oneite hardware installation and packaging material remova? erremote guidance for hardware installation ¹	·	Rémote puidance or onsite	Onaite
Sector Sector	Install and configure system software		Rentatio	consta
	Install sapport software and connect with Deli Technologies			•
	Project documentation with knowledge transfer		•	
	Deployment vertication	4	•	
	Configuration data transfer to Dell Technologies technical support	e: 1 fi	•	(O)
Post-deployment	30-days of post-deployment configuration assistance			•
	Training credits for Dell Technologies Education Services			1 • 1
Online oversight	Online collaborative environment in TechDiract for planning, managing and backing delivery ²		- •	

¹ Pockaging removal included with unsite hardware installation final-oded with ProDeploy or PhoDeploy Plus, Not included with Basic Deployment

Figure 45. ProDeploy Infrastructure Suite - Field services

Dell ProDeploy Plus for Infrastructure

From beginning to end, ProDeploy Plus provides the skill and scale that is must successfully perform demanding deployments in today's complex IT environments. Certified Dell experts start with extensive environmental assessments and detailed migration



planning and recommendations. Software installation includes set up of our enterprise connectivity solution (secure connect gateway) and OpenManage system management utilities.

Postdeployment configuration assistance, testing, and product orientation services are also available.

Dell ProDeploy for Infrastructure

ProDeploy provides full-service installation and configuration of both server hardware and system software by certified deployment engineers including set up of leading operating systems and hypervisors as well our enterprise connectivity solution (secure connect gateway) and OpenManage system management utilities. To prepare for the deployment, we conduct a site readiness review and implementation planning exercise. System testing, validation, and full project documentation with knowledge transfer complete the process.

Dell Basic Deployment

Basic Deployment delivers worry-free professional installation by experienced technicians who know Dell servers inside and out.

Additional Deployment Services

You can tailor the ProDeploy Infrastructure Suite offer to meet your customer's unique needs by leveraging "Additional Deployment Time." ADT will cover additional tasks above the normal scope of the standard offers. ADT can be sold for Project Management or Technical Resources and is sold as blocks of four hours remote or eight hours on-site.

Dell ProDeploy for HPC (available in US/Canada only. All other regions use custom)

HPC deployments require specialists that understand that cutting edge is yesterday's news. Dell deploys the world 's fastest systems and understands the nuances that make them perform. ProDeploy for HPC provides:

- Global team of dedicated HPC specialists
- Proven track record, thousands of successful HPC deployments
- Design validation, benchmarking, and product orientation

Learn more at Dell.com/HPC-Services.



ProDeploy Expansion for HPC

*Available as standard SKUs in US & Canada and as custom quote in APJC, EMEA, LATAM



Supplemental Deployment Services

Additional ways to expand scope or deploy for unique scenarios.

Two Host Adder (requires PD/PDP)

Deploying new storage, compute, or networking devices may require interconnection to other servers (also called hosts). The Dell delivery team will set up four hosts per device as part of every ProDeploy service. For example, if the customer is buying two storage arrays the ProDeploy service will automatically include connectivity of four hosts each (4x2=8 total hosts per project since there are two devices). This supplemental "Two Host Adder" service provides for the configuration of additional hosts above what is already provided as part of the ProDeploy service. In many cases, customers can work with us while we set up the included hosts, so they may understand how to do the rest themselves. Always ask the customer how many hosts are being connected and sell the host adder depending on the customer's technology skillset. Note that this service applies to the connectivity of Dell devices not 3rd party devices.

Additional Deployment Services (ADT) – sold with or without PD/PDP

You can expand the scope of a ProDeploy engagement leveraging Additional Deployment Time (ADT). ADT covers additional tasks above the normal deliverables of the ProDeploy offers. ADT can also be used as a standalone service without ProDeploy. SKUs are available for both Project Management and Technical Resource Expertise. SKUs are sold as blocks of four hours remote or eight hours onsite. The delivery team can help in scoping the number of hours required for additional tasks.

Data Migration Services

Migrating data sets is no easy task. Our experts use proven tools and process to streamline data migrations and avoid compromising data. A customer project manager works with our experienced team of experts to create a migration plan. Data migration is part of every technology upgrade, platform change, and shift to the cloud. You can rely on Dell data migration services to perform a seamless transition.



Residency Services

Certified technical professionals act like an extension of your IT staff to enhance internal capabilities and resources and help you realize faster adoption and maximized ROI of new technology. Residency Services help customers transition to new capabilities quickly by leveraging specific technology skill sets. Residency experts can provide post implementation management and knowledge transfer that is related to a new technology acquisition or day-to-day operational management of the IT infrastructure.

- Global experts available to serve in-person (onsite) or virtual (remote)
- Engagements starting at 2 weeks with flexibility to adjust
- Residency is available for project management needs, and many different technology skills sets such as: Server, storage, Gen AI, networking, security, multi-cloud, data mgmt., and modern workforce application residents

Unique Deployment Scenarios

Custom Deployment Services

When a deployment is beyond the scope of the ProDeploy Infrastructure Suite, you can turn to the custom deployment services team to address complex implementation scenarios and customer unique requirements. The Dell custom deployment team is staffed with solution architects who will assist with customer scoping calls to define the project and develop the statement of work. Custom services can handle a wide range of deployments that can be performed in the factory or onsite. All custom engagement services are requested through SFDC.

ProDeploy FLEX

ProDeploy Flex is a modular service and a powerful tool for you to attach more services and improve revenue and margins. The ProDeploy Flex modular offer allows sales teams to build and better tailor services by mixing factory and field delivery options. You can also select special deployment scenarios without going to the custom order desk. FLEX is ideal for unique deployments where ProDeploy or ProDeploy Plus are not an adequate answer to the customer needs. Key features of ProDeploy FLEX :

- Build deployment quotes using modular, selectable features for both hardware and software.
- The system automatically scales pricing based on volume.
- Ideal for customers who require NativeEdge Orchestrator or edge deployments.
- Ability to add deployment services to third-party networking devices.

Deployment of HPC

High-Performance Computing (HPC) implementations require specialists that understand advanced feature sets. Dell deploys the world 's fastest systems and understands the nuances that make them perform. HPC deployments are most often scoped as custom service engagements, however we can do smaller HPC clusters under 300 nodes using a standard ProDeploy SKU. Any standard SKU for HPC deployment will be sold as one base SKU per cluster (ProDeploy for HPC Base) along with one ProDeploy for HPC Add-on for each device in the cluster (server nodes and switches).

Scope of ProDeploy for HPC:

(i) NOTE: Available as standard SKUs in US and Canada. Custom Service would be required for all other regions.





(regardless of cluster size)

Figure 47. Standard deliverables of ProDeploy for HPC

HPC Add-on for Nodes

- Rack & Stack Server Nodes
- Professionally labeled cabling
- BIOS configured for HPC
- OS installed
- Per node
- Tied & Non-Tied Add-on SKUs
- 1 SKU/asset
- If over 300 nodes use custom quote

Build HPC solutions for your unique requirements Choose ProDeploy for HPC or Custom deploy

ProDeploy service includes configuration of most OS, cluster mgmt., networking and benchmarking



Notes related to networking above: Omni-Path is no longer an Intel Product, but is now distributed by a company called Cornelis, and Mellanox was purchased by Nvidia, and now goes by Nvidia Networking.

Figure 48. Visual view of HPC deployment options to include hardware and software

DAY 2 – Automation Services with Ansible

Dell solutions are built as "automation ready" with integrated APIs (Application Programming Interfaces) to allow customers to programmatically call actions on the product through code. Although Dell has published Anisble automation use cases, some customers need additional assistance with GitOps. By the end of the service, the customer will have the foundational

Appendix D: Service and support 77



components required to accelerate automation and understand how the programming works together: Day 1 and Day 2 use case value automation scripts (ansible modules), CI/CD tool (Jenkins), and Version control (Git).

Dell Technologies Consulting Services

Our expert consultants help customers transform faster, and quickly achieve business outcomes for the high value workloads Dell PowerEdge systems can handle. From strategy to full-scale implementation, Dell Technologies Consulting can help determine how to perform IT, workforce, or application transformation. We use prescriptive approaches and proven methodologies that are combined with portfolio and partner ecosystem of Dell Technologies to help achieve real business outcomes. From multi cloud, applications, DevOps, and infrastructure transformations, to business resiliency, data center modernization, analytics, workforce collaboration, and user experiences-we are here to help.

Dell Managed Services

Some customers prefer Dell to manage the complexity and risk of daily IT operations, Dell Managed Services utilizes proactive, Al enabled delivery operations and modern automation to help customers realize desired business outcomes from their infrastructure investments. With these technologies, our experts run, update and fine-tune customer environments aligned with service levels, while providing environment-wide and down-to-the-device visibility. There are two types of managed service offers. First the outsourcing model or CAPEX model where Dell manages the customer owned assets using our people and tools. The second is the as-a-Service model or OPEX model called APEX. In this service, Dell owns all technology and all the management of it. Many customers will have a blend of the two management types depending on the goals of the organization.

Managed	CAPEX model	 APEX	as-a-Service or OPEX model
We manage using our pe • Managed o • Technology • End-user (I • Service des • Cloud Man • Office365 o	your technology ople and tools. ¹ letection and response* / Infrastructure PC/desktop) sk operations aged (Pub/Private) or Microsoft Endpoint	We own can off- • AP • AP ela • AP pay	n all technology so you load all IT decisions. EX Cloud Services EX Flex on Demand stic capacity EX Data Center Utility y-per-use model

* Managed detection and response covers the security monitoring of laptops, servers, & virtual servers. Min. 50 devices combined. No Networking or Storage-only systems [SAN/NAS]. Available in 32 countries. Details here

Figure 49. Dell Managed Services

Managed Detection and Response (MDR)

Dell Technologies Managed Detection and Response (MDR) is powered by Secureworks Taegis XDR software platform. MDR is a managed service that secures the customer's IT environment against malicious actors and provides remediation if and when a threat is identified. When a customer purchases MDR, they will receive the following features from our team:

- Dell badge resources
- Agent rollout assistance to help deploy the Secureworks Endpoint Agent
- 24x7 threat detection & investigation
- Up to 40hrs per quarter of response and active remediation activities
- If the customer experiences a breach, we will provide up to 40hrs per year of Cyber incident response initiation
- Quarterly reviews with the customer to review the data



Dell Technologies Education Services

Build the IT skills required to influence the transformational outcomes of the business. Enable talent and empower teams with the right skills to lead and perform transformational strategy that drives competitive advantage. Leverage the training and certification required for real transformation.

Dell Technologies Education Services offers PowerEdge server training and certifications that are designed to help customers achieve more from their hardware investment. The curriculum delivers the information and the practical, firsthand skills that their team must confidently install, configure, manage, and troubleshoot Dell servers.

To learn more or register for a class today, see Education.Dell.com.

Appendix D: Service and support 79

=			⊕ Q.×
lutos Intel®	Processadores	Intel®	/ Intel [®] Xeon [®] Processors
	Xeon Processador Intel® Xeon® S cacho do 37.5 M 2.00 GHz	Silver 4416+	
	Intel Processador Intel® Xeon® S Xeon cacho do 27.5 M 2.00 CU	Silver 4416+	
	cache de 37,5 M, 2,00 GHZ		
	Adicionar para comparar		
	Especificações		
	Baixe as especificações↓		
	Essenciais		
	Coleção de produtos	Processadores escaláveis Intel® Xeon® da 4ª Geração	
	Codinome	Produtos com denominação anterior Sapphire Rapids	
	Segmento vertical	Server	
	Número do processador ③	4416+	
	Litografia ③	Intel 7	
	Preço recomendado para o cliente 🕲	\$1176.00-\$1186.00	
	Especificações da CPU		
	Número de núcleos ③	20	
	Total de threads ③	40	
	Frequência turbo max ③	3.90 GHz	
	Frequência base do processador 🏵	2.00 GHz	
	Cache ③	37.5 MB	
	Velocidade do Intel® UPI	16 GT/s	
	№ de links de UPI ③	2	
		16E M	



Status	Launched
Data de introdução ③	Q1'23
Status de manutenção ③	Baseline Servicing
Opções integradas disponíveis ③	Sim
Condições de uso 🕥	Server/Enterprise
Especificações de memória	
Tamanho máximo de memória (de acordo com o tipo de memória) ③	4 TB
Tipos de memória ③	Up to DDR5 4000 MT/s 1DPC and 2DPC
№ máximo de canais de memória ③	8
Compatibilidade com memória ECC $^{^{\dagger}}$ ()	Sim
Opções de expansão	
Escalabilidade	2S
Revisão de PCI Express ③	5
№ máximo de linhas PCI Express ③	80
Especificações de encapsulamento	
Soquetes suportados ③	FCLGA4677
Transportadora de pacotes	E1B
DTS Max	94 °C
T _{CASE} ③	82
Tamanho do pacote	77.5mm x 56.5mm
Atualizações disponíveis do Intel® On Den	nand

Activation Model Products	QAT	DLB	DSA	IAA
Communications & Storage Suite 2	2	2		
SGX512				
•				Þ
Tecnologias avançadas				
Ativação do recurso Intel® On Deman	d 🕲 🛛	Sim		
Intel® QuickAssist Technology (QAT)		1 default devices		
Intel® Dynamic Load Balancer (DLB)		1 default devices		
Intel® Data Streaming Accelerator (DS	A)	1 default devices		
Intel® In-memory Analytics Accelerate	or (IAA)	1 default devices		
Intel® Advanced Matrix Extensions (Al	MX)	Sim		

Inserido ao protocolo **22.951.206-4** por: **Pedro Henrique Golin Linhares** em: 31/03/2025 12:49. A autenticidade deste documento pode ser validada no endereço: https://www.eprotocolo.pr.gov.br/spiweb/validarDocumento com o código: **692a5ef82beb173ed33dad4ee6505546**.



Intel® Deep Learning Boost (Intel® DL Boost) ⑦	Sim
Intel® Resource Director Technology (Intel® RDT) ⑦	Sim
Tecnologia Intel® Speed Shift ③	Sim
Tecnologia Intel® Turbo Boost ‡ 🏵	2.0
Tecnologia Hyper-Threading Intel® ‡ 🍞	Sim
Intel® TSX-NI 🗿	Sim
Intel® 64 [‡]	Sim
Extensões do conjunto de instruções 🕥	Intel® AMX, Intel® SSE4.2, Intel® AVX, Intel® AVX2, Intel® AVX-512
№ de unidades de FMA de AVX-512 🍞	2
Segurança e confiabilidade	
Intel® Software Guard Extensions (Intel®SGX) ⑦	Yes with Intel® SPS
Tamanho máximo do cache de página de enclave (EPC) para Intel® SGX	64 GB
Intel® Crypto Acceleration ③	Sim
Aceleração de software Intel® QuickAssist	Sim
Suporte para Resiliência de firmware de plataforma Intel®	Sim
Intel® Control-Flow Enforcement Technology ⑦	Sim
Intel® Total Memory Encryption ③	Sim
Novas instruções Intel® AES 🧿	Sim
Intel® OS Guard	Sim
Intel® Trusted Execution Technology $^{^{\dagger}}$ $\textcircled{3}$	Sim
Bit de desativação de execução ‡ 🏵	Sim
Intel® Boot Guard ③	Sim
Controle de Execução baseado em Modo (MBEC — Mode-based Execute Control) ③	Sim
Tecnologia de virtualização Intel® (VT-x) † 🏵	Sim
Tecnologia de virtualização Intel® para E/S dirigida (VT-d) † 🍞	Sim
Intel® VT-x com Tabelas de páginas estendidas (EPT) [†] ③	Sim

Todas as informações fornecidas estão sujeitas a alterações a qualquer momento, sem aviso prévio. A Intel pode alterar o ciclo de vida da fabricação, as especificações e as descrições dos produtos a qualquer momento, sem aviso prévio. As informações aqui contidas são fornecidas "no estado em que se encontram" e a Intel não atribui qualquer declaração ou garantias relacionadas à precisão das informações, nem sobre os recursos dos produtos, disponibilidade, funcionalidade ou compatibilidade dos produtos listados. Para obter mais informações sobre os produtos ou sistemas, entre em contato com o fornecedor do sistema.

Inserido ao protocolo **22.951.206-4** por: **Pedro Henrique Golin Linhares** em: 31/03/2025 12:49. A autenticidade deste documento pode ser validada no endereço: https://www.eprotocolo.pr.gov.br/spiweb/validarDocumento com o código: 692a5ef82beb173ed33dad4ee6505546.



Intel classifications are for general, educational and planning purposes only and consist of Export Control Classification Numbers (ECCN) and Harmonized Tariff Schedule (HTS) numbers. Any use made of Intel classifications are without recourse to Intel and shall not be construed as a representation or warranty regarding the proper ECCN or HTS. Your company as an importer and/or exporter is responsible for determining the correct classification of your transaction.

Consulte a Ficha técnica para obter definições formais de propriedades e recursos de produtos.

‡ Este recurso pode não estar disponível em todos os sistemas de computação. Verifique com o fornecedor do sistema para determinar se seu sistema oferece este recurso ou consulte as especificações de seu sistema (motherboard, processador, chipset, alimentação, HDD, controle gráfico, memória, BIOS, drivers, monitor de máquina virtual [VMM], software de plataforma e/ou sistema operacional) para saber sobre a compatibilidade do recurso. A funcionalidade, o desempenho e outros benefícios deste recurso podem variar, dependendo das configurações do sistema.

Os números dos processadores Intel não são indicação de desempenho. Os números dos processadores diferenciam recursos dentro de cada família de processador, e não entre famílias diferentes de processadores. Consulte https://www.intel.com.br/content/www/br/pt/processors/processor-numbers.html

para obter mais detalhes.

O RCP (Recommended Customer Price, preço recomendado para o cliente) é o guia de preços somente para produtos Intel. Os preços são para clientes diretos da Intel, representam geralmente as quantidades de compra de 1.000 unidades, e estão sujeitos a alterações sem aviso prévio. Os preços podem variar para outros tipos de pacotes e quantidades de envio. Na venda por atacado, o preço corresponde à unidade. Listar os índices RCP não constituí uma oferta oficial da Intel.

O TDP máximo e do sistema se baseiam nos piores casos. O TDP real pode ser inferior, se nem todas as E/Ss para chipsets forem utilizadas.

SKUs "anunciados" ainda não estão disponíveis. Favor consultar a data de lançamento para a disponibilidade no mercado.

Frequência máxima de turbo refere-se à frequência máxima do processador de núcleo único que pode ser atingida com a Tecnologia Intel[®] Turbo Boost. Mais informações estão disponíveis no site https://www.intel.com/content/www/br/pt/architecture-andtechnology/turbo-boost/turbo-boost-technology.html

Consulte https://www.intel.com.br/content/www/br/pt/architecture-and-technology/hyper-threading/hyper-threading-technology.html? wapkw=hyper+threading

para obter mais informações, incluindo detalhes sobre quais processadores são compatíveis com a Tecnologia Hyper-Threading Intel[®].

Os processadores compatíveis com a computação de 64 bits na arquitetura Intel® requerem BIOS habilitados para arquitetura Intel 64.

Alguns produtos suportam as novas instruções AES com uma atualização da Configuração do processador, em particular, i7-2630QM/i7-2635QM, i7-2670QM/i7-2675QM, i5-2430M/i5-2435M, i5-2410M/i5-2415M. Favor entrar em contato com o OEM para o BIOS que inclui a mais recente atualização da Configuração do processador.

Informações sobre a empresa

Nosso compromisso

Inclusão

Relações com investidores

Fale conosco

Sala de imprensa

Mapa do site

Empregos

f

in

© Intel Corporation Termos de uso *Marcas comerciais Cookies Privacidade Transparência da cadeia de fornecimento

As tecnologias Intel[®] podem exigir ativação de hardware, software específico ou de serviços. // Nenhum produto ou componente pode ser totalmente seguro. // Os seus custos e resultados podem variar. // O desempenho varia de acordo com o uso, a configuração e outros fatores. // Veja nossos Avisos e isenções de responsabilidade legais completos

 \mathbb{X}

. // A Intel está comprometida em respeitar os direitos humanos e evitar cumplicidade com abusos de direitos humanos. Consulte Princípios Globais de Direitos Humanos da Intel. Os produtos e software da Intel são destinados a serem utilizados apenas em aplicações que não causem ou

contribuam com a violação de um direito humano reconhecido internacionalmente.

r para o conteúdo principal



intel.

Inserido ao protocolo **22.951.206-4** por: **Pedro Henrique Golin Linhares** em: 31/03/2025 12:49. A autenticidade deste documento pode ser validada no endereço: https://www.eprotocolo.pr.gov.br/spiweb/validarDocumento com o código: **692a5ef82beb173ed33dad4ee6505546**.



Cyber Resilient Security in Dell PowerEdge Servers

October 2023 H19738.1

White Paper

Abstract

This white paper highlights the Dell PowerEdge Cyber Resilient Architecture and describes the server life cycle for implementing zero trust principles for your infrastructure. Dell PowerEdge security controls provide a comprehensive security solution that ensures resiliency while enforcing a zero-trust posture.

D

L

Technologies

Inserido ao protocolo **22.951.206-4** por: **Pedro Henrique Golin Linhares** em: 31/03/2025 12:49. A autenticidade deste documento pode ser validada no endereço: https://www.eprotocolo.pr.gov.br/spiweb/validarDocumento com o código: fc1c44dfe0686d5547ffd691199be699.



Copyright

The information in this publication is provided as is. Dell Inc. makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Use, copying, and distribution of any software described in this publication requires an applicable software license.

Copyright © 2023 Dell Inc. or its subsidiaries. Published in the USA October 2023 H19738.1.

Dell Inc. believes the information in this document is accurate as of its publication date. The information is subject to change without notice.



Contents

Contents

Executive summary	.5
Overview	.5
Revisions	.5
We value your feedback	.5
Introduction	.6
Digital infrastructure complexity	.6
Sophistication and complexity of threats	.6
Regulatory landscape and internal mandates	.6
Zero Trust strategy for the modern world	.6
Core principles of Zero Trust	.7
Seven pillars of Zero Trust	.8
The Dell advantage	.9
Security journey across the server life cycle1	0
First stage – Choosing the server1	0
Challenges1	0
PowerEdge solutions	0
Dell Secure Development Lifecycle	1
Compliance advantage	1
Rapid response to new vulnerabilities	2
Bug Bounty Program	2
Solutions covering threat vectors for every layer of the server1	3
Second stage – Supply chain security1	4
Challenges1	4
PowerEdge solutions1	4
End-to-end supply chain assurance1	4
Secured Component Verification1	6
Software Bill of Materials1	6
Third stage – Efficient deployment and configuration at-scale	7
Challenges	7
PowerEdge solution	7
System integrity	7
Hardware security	20
Protecting data at rest	21
Protecting Data in Flight	23
Protecting data in use	26
Identity Access Management	28
Capabilities and automation for efficient at-scale deployment	30



Contents

Fourth Stage – Security management and monitoring	
Challenges	
PowerEdge Solutions	31
Visibility, logging, and alerts	31
SELinux framework	
Real-time detection – BIOS Live scan	
Silicon-based Root of Trust	34
Automated and manual recovery	
Updating	
Restoring server configuration after hardware servicing	
CloudIQ	
Managed detection and response services	39
Fifth stage – Secure decommissioning and repurposing	40
Challenge	40
PowerEdge solutions	40
Secure Erase	41
Secure erase – physical disk	41
Data sanitization and destruction services	42
Summary	42
References	44
Dell Technologies documentation	44



Executive summary

Overview

Revisions

The Dell Technologies approach to security is intrinsic in nature – it is built-in, not boltedon later, and it is integrated into every step of Dell's Secure Development Lifecycle. We continuously strive to evolve our Dell PowerEdge security controls, features, and solutions to meet the ever-changing threat landscape and to help customers accelerate Zero Trust adoption.

Securing your infrastructure is not a one-time investment, but a mindset and overall approach. This white paper uses this journey perspective to describe the PowerEdge advantages across the server life cycle. In each of the life cycle phases, from deployment through maintenance to decommissioning, we highlight how Dell's PowerEdge Cyber Resilient Architecture security features work together to provide both resiliency and a zero-trust approach. The Dell Remote Access Controller (iDRAC9) enables many of these features.

We continue to anchor security with a Silicon-based Root-of-Trust (RoT). Since the previous PowerEdge cyber resilient security whitepaper, many new features have been added that span from access control to data encryption to supply chain assurance. All features make extensive use of intelligence and automation to help you stay ahead of the threat curve, and to enable the scaling demanded by ever-expanding use models. Dell's Cyber Resilient Architecture, enhanced over many years, is the foundation for the critical elements of a Zero Trust environment.

Date	Part number/ revision	Description
November 2022	H19738	Initial release
October 2023	H19738.1	Updated to include the 16 th generation of PowerEdge servers

We value yourDell Technologies and the authors of this document welcome your feedback on this
document. Contact the Dell Technologies team by email.

Authors: Deepak Rangaraj, Kim Kinahan

Contributors: Marshal Savage

Note: For links to other documentation for this topic, the <u>Dell Technologies Info Hub for</u> <u>PowerEdge</u>.



Introduction

Digital infrastructure complexity	Modern IT environments have changed drastically in the past few years with servers being deployed in various use cases such as on-premises, multicloud, Edge, Telco, and so on. The server platforms are becoming more complex with an ever-increasing number of components that require firmware for configuration and management. We are generating data at a speed and volume higher than ever and this data is often generated and stored at numerous locations distributed geographically. This increasing complexity necessitates effective management of security controls to mitigate the expansion of the attack surface.
Sophistication and complexity of threats	The Dell Technologies Digital Transformation Index found that data privacy and cybersecurity concerns are the leading barriers to digital transformation ¹ . The complexity, sophistication, and frequency of cyberattacks are increasing, and the damage caused by attacks is becoming more costly. Complicating matters further, today's threat actors are taking advantage of technological advancements, such as AI and a lower cost of entry. Malicious actors are continuously searching for vulnerabilities to exploit. With the assistance of advanced AI systems, they can carry out nefarious activities at an unprecedented scale and manipulate systems in innovative and harmful ways beyond human capabilities. Global damages related to cybercrime are predicted to reach \$10.5 trillion by 2025. ²
Regulatory landscape and internal mandates	As global threats increase, governments worldwide are developing regulatory guidance in response to cyber threats. As a result, private institutions are creating stronger policies and mandates to mitigate advanced persistent threats. Also, there are more mandates for security requirements that are needed to work with the government. These requirements impact suppliers, vendors, and any organization that partners with the government. In addition, regulations are carrying over to the critical infrastructure sectors, such as healthcare, transport, and finance. Outside of government regulations, many customers want to harden their infrastructure and are developing their own internal mandates or security policies.
Zero Trust strategy for the modern world	The increasing infrastructure complexity and threat landscape are driving a critical need to secure not only infrastructure hardware, but also the firmware and the supply chain itself. When applying a Zero Trust strategy, customers focus on business controls, the control plane, and applications and data. However, there is a critical need to secure what is below these items, including infrastructure hardware and firmware, supply chain for the infrastructure, and design and processes used to build the infrastructure. Zero Trust principles must be applied to all these aspects for more comprehensive cyber resilience.
	Dell Technologies has security built into our industry-leading servers, storage, HCI, and data protection appliances to help protect data wherever it is stored, managed, or used. As a foundation for securing our PowerEdge server products, they are foremost cyber resilient – capable of anticipating, withstanding, and recovering from cyber threats.
	¹ Dell Technologies 2020 Digital Transformation Index
	² www.cybersecurityventures.com /cybercrime-damage-costs-10-trillion-by-2025/

951

Simultaneously, the process of designing PowerEdge security controls and tools incorporates Zero Trust principles. By anticipating how customers want to use these capabilities while they are setting up their Zero Trust deployments, we have adapted our approach. We made it easier to work with us no matter where customers are on their journey towards Zero Trust adoption.



Figure 1. Infrastructure security and Zero Trust

The Dell infrastructure security approach is integral to Dell products and accelerates Zero Trust adoption because it is:

- Designed and built using Zero Trust principles
- Provides Zero Trust capabilities and features
- Provides common, consistent behavior

Core principles of Zero Trust

The tenets of Zero Trust architecture are built on a set of principles that presume that the *network is always vulnerable to compromise* and sets out to safeguard access to critical data and resources. Unlike trust-then-verify frameworks, a zero-trust approach eliminates implicit trust. Every user, device, and application must be continuously *authenticated and explicitly authorized* based on a range of factors such as identity, device status, location, and behavior.

Identity plays a crucial role in Zero Trust. Identification pertains not only to people but to applications, communication paths, network devices, and the data itself. When an IT asset is identified, authenticated, and explicitly authorized, the principle of *least privilege* is applied. This approach ensures that only authorized entities are given a minimum level of access required to perform their specific task. The data-centric security model constantly limits access while also looking for anomalous or malicious activity. The Zero Trust approach reduces the granularity of validation at key intersections for verified trust, optimizing least privilege without impacting workload efficiency. The goal is to deter attacks and reject them at point of entry. However, if a breach occurs, the amount of damage is minimized, along with an enhanced ability to detect and remediate immediately.

Adopting the Zero Trust mindset and using Zero Trust principles enables systems administrators to control how users, processes, and devices engage with data. These



Introduction

principles can prevent the abuse of compromised user credentials, remote exploitation, insider threats, and even mitigate the effects of malicious supply chain activity.

Seven pillars of Zero Trust The Zero Trust model, defined by the National Institute of Standards and Technology (NIST)³, identifies seven interrelated pillars that work together to provide a comprehensive and holistic approach to infrastructure and data security. Each pillar represents a specific functional or key focus area for implementation of Zero Trust security controls. When combined, these seven pillars provide a multifaceted, layered, and integrated security framework.

Dell's Zero Trust approach integrates a broad set of security controls and automation capabilities for the management of the infrastructure and the applications running on it. The following table highlights Dell's capabilities across the seven pillars as outlined by NIST:

Zero Trust pillar	NIST description ³	PowerEdge highlights
User	 User identification, authentication, and access control: Only validated and authorized users can access data and resources. The principle of <i>least privilege</i> is applied where users are granted the minimum level of access required to perform their specific tasks. 	 Identity and Access Management Multi-Factor Authentication—RSA Secure ID Active Directory or LDAP integration with single sign-on (SSO) support Role-based access control and auditing
Device	 Monitoring and enforcement of device health, compliance, and device posture assessment: Monitoring - looking for anomalies and suspicious read/write activity. Health – confirming the latest version of the firmware. All devices are identified, inventoried, authorized, authenticated, and updated. 	 Silicon Root-of-Trust (ROT) with complementary Intel Boot Guard and AMD Platform Secure Boot (AMD PSB) Secure supply chain with Secured Component Verification (SCV) Chassis locks and intrusion detection Dynamic USB port enable/disable Trusted Platform Module (TPM) Device attestation with SPDM (Security Protocol Data Model from DMTF)
Data	Ensure data transparency and visibility by using enterprise infrastructure, applications, standards, solid end-to-end encryption, and data tagging.	 Data-at-rest protection: Drive encryption with local (LKM) or Secure Enterprise Key Management (SEKM) with direct-attached NVMe drive support Baseboard Management Controller (BMC)-based Local Key Management (iLKM) Data-in-use protection:

Table 1. Dell Zero Trust approach across the seven pillars

³ National Institute of Standards and Technology (NIST) Special Publication (SP) 800-207, Zero Trust Architecture

8 Cyber Resilient Security in Dell PowerEdge Servers

Inserido ao protocolo **22.951.206-4** por: **Pedro Henrique Golin Linhares** em: 31/03/2025 12:49. A autenticidade deste documento pode ser validada no endereço: https://www.eprotocolo.pr.gov.br/spiweb/validarDocumento com o código: fc1c44dfe0686d5547ffd691199be699.



Introduction

Zero Trust pillar	NIST description ³	PowerEdge highlights
		 Confidential compute—Intel SGX, Intel MKTME, AMD Secure Memory Encryption (SME) AMD SME, AMD Secure Encrypted Virtualization (SEV) SEV, AMD persistent memory encryption
Application and Workload	Secure applications and workloads, and protect containers and VMs.	 Secure Development Lifecycle Cryptographically signed BIOS and firmware updates Secure end-to-end boot and Unified Extensible Firmware Interface (UEFI) boot capabilities Drift detection Rapid Response and mitigations for CVEs
Network and Environment	Encrypt, monitor, and analyze network. Logically and physically segment, isolate, and control the network and the environment (on- premises and off-premises) using granular access and policy restrictions.	 Dedicated BMC (iDRAC) network module SSH/TLS communication options TLS 1.3 support DPU/SmartNIC
Visibility and Analytics	Monitor activities and behaviors across the infrastructure (user, device, data, network, and application) to identify patterns and anomalies. Use analytics to detect and respond to security threats.	 Persistent event logging and auditing Real-time and boot time firmware scanning Security alerts CloudIQ
Automation and Orchestration	Automate manual security and other applicable processes to take policy-based actions across the enterprise with speed and at scale.	 OpenManage Enterprise drift detection Firmware rollback Automatic BIOS and operating system recovery Centralized updates Automatic SSL certificate renewal

The Dell advantage

Security is in our DNA, and we are committed to making our products secure by design and secure by default. PowerEdge servers are built using Zero Trust principles internally and have capabilities that enable customers to set up a Zero Trust IT environment and operations. Our products also strive to provide common and consistent behavior and controls across our portfolio.

PowerEdge servers with iDRAC9 have an integrated immutable silicon-based platform Root-of-Trust (RoT) that is used to establish a verified chain of trust that extends throughout the server life cycle, from deployment through maintenance to decommissioning. This RoT combined with security controls and comprehensive management tools provides robust layers of security across the PowerEdge hardware and firmware.



First stage – Choosing the server

These PowerEdge capabilities not only ensure cyber resiliency to protect, detect, and recover from attacks, they also maintain a locked-down posture for a Zero Trust approach of least privilege. Least privilege ensures that users and devices are only given access to what they need to perform their tasks. Our goal is to make Zero Trust a reality for our customers and accelerate its speed of adoption.

Security journey across the server life cycle

Implementing security across your infrastructure requires a series of ongoing efforts and measures to protect underlying systems, networks, and resources. Transitioning to a mature security model is not a one-time investment and cannot be accomplished overnight. It is an ongoing journey and approach of implementing stringent security policies. As your Zero Trust implementation matures over time, enhanced visibility and extensive controls allow you to keep pace with the threat landscape. We have organized this journey into five stages.

- **Stage 1: Choosing the server**—Customers want assurance that security is a top priority and is built into every aspect of design.
- Stage 2: Ensuring supply chain security—Customers face the real risk of malicious offenders replacing original components with counterfeits, implants, or malware.
- Stage 3: Efficient deployment and configuration at scale—How a server is deployed has a direct impact on its performance, stability, and security. Proper planning and configuration are essential to ensure that the server is set up correctly and that all necessary components are in place.
- Stage 4: Security management and monitoring— Because attacks happen quickly, faster than a human can detect, customers must proactively monitor their environment to take quick action. The lack of skills and training can also exacerbate the problem.
- Stage 5: Secure repurposing and decommissioning—Data security is a key consideration when the server is repurposed or retired. IT best-practices recommend removing all data from the server to ensure that confidential information is not inadvertently shared or compromised.

First stage – Choosing the server

Challenges Customers want to be confident that their server is secure in its entirety and does not introduce any vulnerabilities into their environment. They want assurances from server vendors addressing every aspect of design, the entire supply chain, including hardware, software, and firmware. This level of assurance differentiates on-premises infrastructure when compared to cloud service providers who might only offer a black box infrastructure, leaving customers uncertain about the security of the underlying components of their cloud-based system.

PowerEdge
solutionsThe first step of the Zero Trust journey with Dell Technologies starts before you receive
your PowerEdge server. Intrinsic security practices are incorporated into hardware
product design and software or firmware code development. These practices include
processes and policies that ensure security features are implemented at the time of
product inception and continue throughout the development cycle. In essence, security is

10 Cyber Resilient Security in Dell PowerEdge Servers

Inserido ao protocolo **22.951.206-4** por: **Pedro Henrique Golin Linhares** em: 31/03/2025 12:49. A autenticidade deste documento pode ser validada no endereço: https://www.eprotocolo.pr.gov.br/spiweb/validarDocumento com o código: fc1c44dfe0686d5547ffd691199be699.
ls. 955



Dell Secure Development Lifecycle

Delivering the Cyber Resilient Architecture requires security awareness and discipline at each stage of development. The Secure Development Lifecycle (SDL) model is used to make security an integral part of the overall server design process. Some key aspects of this process include:

- Features that are conceived, designed, prototyped, implemented, set into production, deployed, and maintained with security as a key priority
- Server firmware that is designed to obstruct, oppose, and counter the injection of malicious code during all phases of the product development life cycle:
 - Threat modeling and penetration testing provide coverage during the design process.
 - Secure coding practices are applied at each stage of firmware development.
- For critical technologies, external audits that supplement the internal SDL process to ensure that firmware adheres to known security best practices
- On-going testing and evaluation of new potential vulnerabilities using the latest security assessment tools
- Rapid response to critical Common Vulnerabilities and Exposures (CVEs) including recommended remediation measures if warranted



Figure 2. Dell Secure Development Lifecycle

Compliance advantage

Dell Technologies has received the certifications that are needed to comply with major United States Federal and other global governmental requirements as well as industry standards such as from NIST, as described in the following table:



First stage - Choosing the server

Table 2. Certification descriptions

Certification	Description	
Common Criteria	Certain configurations of PowerEdge Servers include components with common criteria certifications. (For example, IDRAC and TPM)	
FIPS 140	Certain configurations of PowerEdge Servers include FIPS 140-certified cryptographic modules (For example, TPM, IDRAC9, Chassis Management Controller (CMC), Self-Encrypting Drives (SEDs), and SSDs)	
IPv6	 PowerEdge Servers are fully USGv6r1 and IPv6 Ready Logo- compliant and certified with IPv6-only capabilities while running operating systems such as Red Hat Enterprise Linux 8.4 and the applicable versions of Windows 2019 or Windows 2022 Server. 	
	 Dell PowerEdge iDRAC9 (with 5.1x firmware) is USGv6r1 and IPv6 Ready Logo-compliant and certified with IPv6-only capabilities. 	

Rapid response to new vulnerabilities

CVEs are newly discovered attack vectors that compromise software and hardware products. Timely responses to CVEs are critical to most companies so they can swiftly assess their exposure and take appropriate action.

Dell Technologies works aggressively to respond quickly to new CVEs in our PowerEdge servers and provide timely information including the following:

- The products that are affected
- Remediation steps
- If needed, availability of updates to address the CVEs

Bug Bounty Program

Dell Technologies recognizes the value of the security research community to broaden visibility into potential vulnerabilities and threats and welcomes the opportunity to collaborate with community members who share this common goal.

Dell's Bug Bounty Program applies to security vulnerabilities identified in Dell-branded or currently supported products.

Solutions covering threat vectors for every layer of the server

There are many threat vectors in today's changing landscape. The following tables summarize the Dell approach to managing critical threats in each of the server layers.

Server platform layers				
Security layer	Threat vector	Dell solution		
Physical server	Server/component tampering or theft of component	 Secured Component Verification (SCV) Chassis Intrusion Detection Secure Enterprise Key Management Intel TME SPDM 		
Firmware and software	Firmware corruptionMalware injection	 Silicon-based Root-of-Trust Intel Boot Guard AMD Secure Root-of-Trust UEFI Secure Boot Customization Cryptographically signed and validated firmware 		
	Software	CVE reportingPatching as required		
Attestation trust features	Server identity spoofing	 TPM Intel TXT Chain of trust 802.1x features SPDM 		
Server management	Rogue configuration and updatesUnauthorized open-port attacks	iDRAC9 Remote attestation		

 Table 3.
 Dell solution for common server platform layer threat vectors



Second stage - Supply chain security

Server environment layers					
Security layer	Threat vector	Dell Technologies solution			
Data	Data breach	 Self-Encrypting Drives (SED) – FIPS or Opal/TCG Secure Enterprise Key Management ISE-only (Instant Secure Erase) drives Secure User Authentication 			
Supply Chain Integrity	Counterfeit componentsMalware threats	 ISO9001 certification for all global server manufacturing sites Secured Component Verification Proof of possession Software Bill of Materials (SBOM) Security measures implemented as part of the Secure Development Lifecycle 			
Supply Chain Security	 Physical security in manufacturing sites Theft and tempering during transport 	 Transported Asset Protection Association (TAPA) facility security requirements Customs-Trade Partnership Against Terrorism (C-TPAT) Secured Component Verification 			

Table 4. Dell solution for common server environment layer threat vectors

Second stage – Supply chain security

Challenges

Modern server platforms are becoming more complex with hundreds of components that require firmware for configuration and management. As a result, the server supply chain is also becoming increasingly complex with hundreds of third-party vendors supplying components as well as the use of open-source software. This complex supply chain contributes to an increase in the attack surface available for threat actors if not managed properly. Customers face the risk of vulnerabilities and threats being introduced into their environments if the supply chain integrity is not assured. The two main aspects of supply chain integrity include:

- **Maintaining hardware integrity**—Ensuring that there is no product tampering or the insertion of counterfeit components or malicious implants before the product is shipped to customers.
- **Maintaining software integrity**—Ensuring that no malware is inserted in firmware or device drivers before shipping the product to customers and preventing code with known vulnerabilities from being introduced into the environment.

inventory, information, intellectual property, and people. These security measures provide

 PowerEdge solutions
 End-to-end supply chain assurance

 Dell Technologies employs a multifaceted approach to protect its supply chain and delivers solutions that customers can trust in an increased threat environment. Our supply chain security consists of prevention and detection controls that protect physical assets,

14 Cyber Resilient Security in Dell PowerEdge Servers

Inserido ao protocolo **22.951.206-4** por: **Pedro Henrique Golin Linhares** em: 31/03/2025 12:49. A autenticidade deste documento pode ser validada no endereço: https://www.eprotocolo.pr.gov.br/spiweb/validarDocumento com o código: fc1c44dfe0686d5547ffd691199be699.

Second stage – Supply chain security

supply chain assurance and ensure integrity by reducing opportunities for the malicious or negligent introduction of malware and counterfeit components into the supply chain.

Dell supply chain controls span supplier selection, sourcing, production processes, and governance through auditing and testing. When a supplier has been selected, the new product introduction process verifies that all materials used during all build stages are sourced from the approved vendor list and match the bill of materials as appropriate. Material inspections during production help identify components that are mismarked, deviate from normal performance parameters, or contain an incorrect electronic identifier.

Parts are procured directly from the Original Design Manufacturer (ODM) or Original Component Manufacturer (OCM) when possible. The material inspection that occurs during the new product introduction process provides multiple opportunities to identify counterfeit or corrupted components that might have entered the supply chain.

Additionally, Dell Technologies maintains ISO 9001 certification for all global manufacturing sites. Strict adherence to these processes and controls helps minimize the risk of counterfeit components being embedded among the Dell products or malware being inserted into firmware or device drivers. As part of SDL, Dell Technologies has several long-standing, key practices that establish and maintain security in manufacturing facilities and logistical networks.

Facilities used to design, build, customize, or fulfill Dell products must demonstrate compliance with several internationally recognized physical security standards such as those defined by the Transported Asset Protection Association (TAPA), American Society for Industrial Security (ASIS), International Standards Organization (ISO), and the Business Alliance for Secure Commerce (BASC).

Protective measures have also been put in place to guard products against theft and tampering during transport as part of an industry-leading logistics program. This program provides a continuously staffed command center to monitor select inbound and outbound shipments across the globe to ensure that shipments make it from one destination to another without disruption.

Dell Technologies audits suppliers and facilities, addressing various factors, including the use of digital closed-circuit TV cameras, access control systems, intrusion detection, and guard service protocols. Other controls are applied to protect Dell cargo during the shipping and logistics process, including tamper-evident packaging, cargo locks and seals, and threat intelligence monitoring of key freight lanes. Internet of Things (IoT) tracking devices are also deployed on select shipments to enable real-time telemetry data monitoring to escalate any security noncompliance events observed during transit.

Dell Technologies also maintains certifications in multiple secure trade and commerce programs such as Tier 3 status with the United States Customs and Border Protection's Customs Trade Partnership Against Terrorism (C-TPAT), Canada's Partners in Protection (PIP), Singapore's Secure Trade Partnership, and Authorized Economic Operator (AEO) status in several other nations. These programs are internationally recognized by member states of the World Customs Organization and demonstrate "best in class" supply chain security standards in the private sector. These programs focus on supplier accountability, security management policies, counter smuggling, trafficking controls, and tamper prevention – all intended to secure trade across international borders.



Second stage - Supply chain security

Supply chain integrity ensures that customers' products are safely delivered and when received, operate as intended. An important feature of supply chain integrity is the development of hardware and software baseline specifications that are preserved securely and later used as a reference to verify that no unauthorized modifications have been made.

Secured Component Verification

Dell Technologies' Secured Component Verification (SCV) for PowerEdge is a supply chain assurance offering that verifies that the server received by a customer matches the configuration that was shipped from the factory. The factory generates a certificate that contains unique component IDs for a specific server. This certificate is stored in a cryptographically secure vault in iDRAC. On receiving the server, the customer runs the SCV application on the host to generate an inventory of the current system, including unique component IDs, and then validates it against the golden factory inventory in the SCV certificate stored in iDRAC.

The SCV application generates a report that identifies any component mismatches from what was installed in the factory. It also verifies the certificate and Chain of Trust along with the Proof of Possession of the SCV Private key for iDRAC. The current implementation supports direct ship customers and does not include Value Added Reseller (VAR) or Part Replacement scenarios.





Software Bill of Materials

As part of Dell's software supply chain security controls, NIST standards, and in alignment with the President's Executive Order (EO) 14028, a Software Bill of Materials (SBOM) is available for a limited number of products across our portfolio. Dell SBOM data adheres to the Software Package Data Exchange (SPDX) standard and is provided in JSON format. SBOM data provides software supply chain transparency and can be used in vulnerability scanning and asset tracking tools on the customer side.

SBOM enables customers to gain a clearer understanding of the software components, versions, licenses, and any open-source software used on the platform. It can facilitate faster detection of known security vulnerabilities in the software components, ultimately enhancing security.

Challenges

Cyber resilient configuration and deployment of the servers is a critical step in the server life cycle. Any mistakes or oversights during this process can result in poor server performance, downtime, or even security breaches. The right set of controls and gates must be in place to ensure system and data integrity throughout the server operation. This configuration and deployment must be performed for hundreds or thousands of servers while ensuring consistency and minimizing any manual errors.

PowerEdge

solution

System integrity

Ensuring system integrity is foundational for securing the server and establishing a lockdown posture for Zero Trust operations. This integrity starts with ensuring that server hardware and components are genuine, and from a trusted and authorized source. Then, the firmware and software must be verified to ensure that a bad actor has not tampered with them. For PowerEdge, this process of ensuring system integrity starts with a siliconbased platform RoT. This RoT anchors the other security controls on the server platform and establishes a chain of trust for cryptographic verification of hardware and software components on the server.

RoT as the anchor for cryptographically verified trusted booting

One of the most critical aspects of server security is ensuring that the boot process can be verified as secure. This process provides a trusted anchor for all subsequent operations such as booting an operating system or updating firmware.

All PowerEdge servers have an immutable, silicon-based RoT burnt into them from the factory. The RoT has one-time programmable, read-only public keys that can be used for cryptographic verification and attestation of integrity.

The BIOS boot process uses Intel Boot Guard technology or AMD PSB technology that cryptographically verifies the BIOS code to be loaded. A verification failure results in a server shutdown and the Lifecycle Controller Log includes a notification. The IT administrator can then initiate the BIOS recovery. If Boot Guard validates successfully, a chain of trust procedure validates the remaining BIOS modules until control is handed off to the operating system or hypervisor.

In addition to Boot Guard's verification mechanism, iDRAC9 4.10.10.10 or higher provides a RoT mechanism to verify the BIOS image at the host boot time. The host can boot only after the BIOS image is successfully validated. iDRAC9 also provides a mechanism to validate the BIOS image at run time, on demand, or at user-scheduled intervals.

Cryptographically verified Trusted Booting

One of the most critical aspects of server security is ensuring that the boot process can be verified as secure. This process provides a trusted anchor for all subsequent operations such as booting an operating system or updating firmware.

PowerEdge servers have used silicon-based security for several generations for features such as iDRAC Credential Vault, an encrypted secure memory in iDRAC for storing sensitive data. The boot process is verified using a silicon-based RoT to meet recommendations in NIST SP 800-147B (BIOS Protection Guidelines for Servers) and NIST SP 800-155 (BIOS Integrity Measurement Guidelines).



Security Protocol Data Model for component attestation

The Distributed Management Task Force (DMTF), of which Dell Technologies is a leading member, defines the Security Protocol Data Model (SPDM). SPDM defines a consistent, open-standard method of communicating in the server to gather information about server components. This component information is protected by encrypting it and using authenticated key exchange to integrity-protect all communications between components. iDRAC SPDM implementation provides visibility into the PERC12 and certain NIC components. As part of the hardware inventory, iDRAC verifies the authenticity and integrity of PERC12 and NIC devices by cryptographically verifying the identity, firmware, and configuration.

TPM Support

PowerEdge servers support two versions of TPM:

- TPM 2.0 FIPS + Common Criteria + TCG certified (Nuvoton)
- TPM 2.0 China (NationZ)

TPM can be used to perform public key cryptographic functions, compute cryptographic hash functions, generate, manage, and securely store keys, and perform attestation. Intel's Trusted Execution Technology (TXT) functionality and Microsoft's Platform Assurance feature in Windows Server 2016 is also supported. TPM can also be used to enable the BitLocker hard drive encryption feature in Windows Server 2012, Windows Server 2016, and Windows Server 2022.

Attestation and remote attestation solutions can use TPM to take measurements at boot time of a server's hardware, hypervisor, BIOS, and operating system, and compare them in a cryptographically secure manner against "golden" or "base" measurements. These measurements are commonly stored outside the TPM on a remote attestation server solution. The TPM PCR measurements stored in the TPM are recalculated on every boot. If they are not identical, the server system might have been compromised and system administrators can disable and disconnect the server either locally or remotely.

Servers can be ordered with or without TPM, but for many operating systems and other security provisions, it is becoming a standard. TPM is enabled through a BIOS option. It is a Plug-In Module solution; the planar has a connector for this plug-in module.

UEFI Secure boot for firmware

PowerEdge servers also support industry-standard UEFI Secure Boot that checks the cryptographic signatures of UEFI drivers and other code loaded before the operating system is running. Secure Boot represents an industry-wide standard for security in the preboot environment. Computer system vendors, expansion card vendors, and operating system providers collaborate on this specification to promote interoperability.

When enabled, UEFI Secure Boot prevents unsigned (that is, untrusted) UEFI device drivers from being loaded, displays an error message, and does not allow the device to function. You must disable Secure Boot to load unsigned device drivers.

In addition, the 14th, 15th, and 16th generations of PowerEdge servers offer the unique flexibility of using a customized boot loader certificate. This certificate is primarily a feature for administrators of Linux environments that want to sign their own operating system boot loaders instead of relying on the default signing certificate provided by Microsoft's UEFI Certificate Authority (CA). Custom certificates can be uploaded by using the preferred

18 Cyber Resilient Security in Dell PowerEdge Servers

Inserido ao protocolo **22.951.206-4** por: **Pedro Henrique Golin Linhares** em: 31/03/2025 12:49. A autenticidade deste documento pode ser validada no endereço: https://www.eprotocolo.pr.gov.br/spiweb/validarDocumento com o código: fc1c44dfe0686d5547ffd691199be699.

iDRAC API to authenticate the customer's specific operating system boot loader. The NSA cites this PowerEdge UEFI customization method for mitigating Grub 2 vulnerabilities in servers.⁴ PowerEdge supports complete customization of Secure Boot, including removal of all industry-standard certificates provided by Microsoft, VMware, or the UEFI CA.

Intel Boot Guard and AMD PSB

Intel Boot Guard and AMD PSB are host processor features that provide strong firmware integrity guarantees, by preventing firmware that is not authorized by the Dell OEM from running on the system. By enabling these features as additional defense-in-depth measures, certain classes of physical attacks are mitigated, such as flash memory replacement or reprogramming, and Time-of-Check-Time-of-Use (TOCTOU) race conditions. All combined, the RoT features in the system make compromise of the trusted computing base (TCB) difficult.

iDRAC/BMC

The Integrated Dell Remote Access Controller (iDRAC) is a Baseboard Management Controller (BMC) that is integrated in Dell PowerEdge servers. iDRAC provides secure and remote server access for many common management functions; administrators can deploy, manage, monitor, update, troubleshoot, and remediate Dell servers from any location without the use of agents and out of band.

iDRAC offers industry-leading security features that adhere to and are certified against well-known NIST standards, Common Criteria, and FIPS 140-2. It is through iDRAC9 that the end user can configure security features to maximize the security posture of the system.

iDRAC credential vault

The iDRAC service processor provides a secure storage memory that protects sensitive data such as iDRAC user credentials and private keys for self-signed SSL certificates. Another example of silicon-based security, this memory is encrypted with a unique immutable root key that is programmed into each iDRAC chip at the time of manufacture. This memory protects against physical attacks where the attacker desolders the chip to gain access to the data.

SELinux framework

SELinux operates at the core kernel level on iDRAC and does not require user input or configuration. SELinux logs security messages when an attack is detected. These log messages indicate when and how an attacker tried to break into the system. These logs are available through SupportAssist to enrolled customers. In future iDRAC releases, these logs will be available in the Lifecycle Controller Logs.

Factory-generated default passwords

By default, all 14th, 15th, and 16th generations of PowerEdge servers ship with a unique, factory-generated iDRAC password to provide additional security. This password is on the pull-out Service Tag on the front of the chassis, next to the server asset label. If you choose this default option, use this password to log in to iDRAC for the first time, rather than using a universal default password. For security purposes, Dell Technologies strongly recommends changing the default password.

⁴ <u>CTR-UEFI-Secure-Boot-Customization-UOO168873-20.PDF (defense.gov)</u>



Hardware security

Hardware security is an integral part of any comprehensive security solution. Some customers want to limit access to external ports, such as USB. Generally, a server chassis does not need to be opened after it has been put into production. At a minimum, customers always want to track and log any such activities. The overall goal is to discourage and limit any physical intrusion.

Chassis intrusion detection and alert

PowerEdge servers provide hardware intrusion detection and logging, with detection working even when no AC power is available. Sensors on the chassis detect when anyone opens or tampers with the chassis, even during transit. Servers that have been opened while in transit generate an entry in the iDRAC Lifecycle log after power is supplied.

Dynamic USB port management

For more security, you can completely disable USB ports. You also have the option of disabling only the USB ports on the front. For example, USB ports can be disabled for production use and then temporarily enabled to grant access to a crash cart for debugging purposes.

iDRAC Direct

iDRAC Direct is a special USB port that is hardwired to the iDRAC service processor for at-the-server debugging and management from the front of the server (cold aisle). It allows you to attach a standard Micro-AB USB cable to this port and the other end of the cable (Type A) to a laptop. A standard web browser can then access the iDRAC UI for extensive debugging and management of the server. If the iDRAC Enterprise license is installed, you can access the operating system desktop using the iDRAC's Virtual Console.

Because you use iDRAC credentials for logging in, iDRAC Direct works as a secure crash cart with the additional advantage of extensive hardware management and service diagnostics. This method is an attractive option for securing physical access to the server in remote locations (host USB ports and VGA outputs can be disabled in this case).

iDRAC Connection View with Geolocation

Connection View enables iDRAC to report the external switches and ports connected to Server I/O.

It is a feature on select networking devices and requires the Link Layer Discovery Protocol (LLDP) to be enabled on connected switches.

Some of the benefits of Connection View enable you to:

- Remotely and quickly check if server I/O modules (LOMs, NDCs, and add-in PCle cards) are connected to the correct switches and ports
- Avoid costly remote dispatch of technicians to remediate wiring errors
- Avoid tracing cables in the server room hot aisles
- Retrieve information for all connections by accessing the UI or by using RACADM commands



Beyond the obvious time and monetary savings, Connection View provides an additional benefit – real-time geolocation of a physical server or VM. Using iDRAC Connection View, administrators can pinpoint a server to see to which switch and port the server is connected. This information helps secure servers from being connected to networks and devices that do not comply with corporate security guidelines or best practices.

Connection View validates the location of the server indirectly by reporting the switch identities to which it is connected. The switch identity helps to determine the geolocation and to assure that the server is not a rogue server in a nonauthorized site, providing another layer of physical security. This information also provides validation that an application or VM has not "crossed" country borders, and that it is running in an approved, secure environment.

Protecting data at rest

Data at rest protection ensures that sensitive data that resides in storage is protected from unauthorized access through encryption and external key management.

Dell Technologies provides:

- Software-based encryption (for example, virtual devices)
- Enterprise key management (for example, SED devices and key management)
- Hardware drive encryption (for example, SED devices)

Whether it is due to internal policies or external compliance, securing data continues to be a high priority for organizations of all sizes.

The 14th, 15th, and 16th generations of PowerEdge servers offer several storage drive options for securing data, as shown in the following figure:





The options start with drives that support Instant Secure Erase (ISE), a new technology to erase user data instantly and securely. The 14th, 15th, and 16th generations of PowerEdge servers offer ISE-capable drives as a default. This white paper describes ISE in more detail later as part of the System Erase feature description.



The next higher security option is Self-Encrypting Drives (SEDs), which offer locking protection that binds the storage drive to the server and RAID card used. This method protects against so-called "smash and grab" theft of drives and the subsequent loss of sensitive user data. When thieves try to use the drive, they do not know the required locking key passphrase and are thwarted from accessing the encrypted drive data. Customers can protect against theft of the entire server by using Secured Enterprise Key Manager (SEKM), which is described in the following section.

NIST FIPS 140-2 certified SEDs offer the highest level of protection. Testing laboratories have accredited drives conforming to this standard. Tamper-resistant stickers are applied to the drive. Dell SED drives have FIPS 140-2 certification by default.

Secured Enterprise Key Manager

OpenManage Secured Enterprise Key Manager (SEKM) delivers a central key management solution to manage data at rest across the organization. It enables you to use an external Key Management Server (KMS) to manage keys that iDRAC can use to lock and unlock storage devices on a PowerEdge server. Using embedded code that is activated with a special license, iDRAC requests that the KMS creates a key for each storage controller, which iDRAC fetches and provides to the storage controller on every host boot so that the storage controller can unlock the SEDs.

The advantages of SEKM over Local Key Management (LKM) include:

- Protection against "Theft of a server" because the keys are not stored on the server and are stored externally and retrieved by connected PowerEdge server nodes (using iDRAC)
- Centralized and scalable key management for encrypted devices with high availability
- Support for the industry-standard Key Management Interoperability Protocol (KMIP), which enables the use of other KMIP-compatible devices
- · Protection of data at rest when drives or the entire server are compromised
- On-drive encryption performance scales with drive count



Figure 2. Secure Enterprise Key Manager (SEKM)

Local Key Management

PowerEdge servers provide the ability to secure SED drives connected to a PERC controller using Local Key Management (LKM).

To ensure user data protection if a drive is stolen, the SED must be locked with a separate key so that it does not decrypt user data unless that key is provided. This key is

referred to as the Key Encryption Key (KEK). The KEK is stored in the PERC, not on an external server.

Set a keyld/passphrase on the PERC controller to which the SED is connected. Then, the PERC controller generates a KEK using the passphrase and uses it to lock the SED. When the drive is powered on, it comes up as a locked SED and encrypts or decrypts user data only when the PERC provides the KEK to unlock it. If a locked drive is stolen, an attacker cannot provide the KEK, and the user data is protected.

The following figure shows the LKM solution:



Figure 3. Local Key Management (LKM)

iLKM

For direct attach NVMe configurations where a PERC RAID controller is not available, iDRAC can be used as the key manager. This solution called OpenManage iLKM, is iDRAC-based and enables key exchange locally. iDRAC acts as a key manager and generates authentication keys that can then be used to secure storage devices. You can transition from iDRAC-based iLKM to iDRAC-based SEKM to upgrade to external key management.

The following figure shows the iLKM solution:



Figure 4. Integrated Local Key Management (iLKM)

Protecting Data in Flight

Data-in-flight protection ensures that data is protected from unauthorized disclosure or interception as it travels across networks or between systems. Sensitive data can be intercepted, stolen, or modified in transit, leading to data breaches, loss of intellectual property, and other security risks.

Distributed and cloud environments where data is constantly moving between systems and across networks, protecting your data through encryption and access controls are important components for data-in-flight protection in a Zero Trust environment.



TLS 1.3

The iDRAC web server uses a TLS/SSL certificate to establish and maintain secure communications with remote clients. Web browsers and command-line utilities, such as RACADM and WS-Man, use this TLS/SSL certificate for server authentication and establishing an encrypted connection.

There are several options available to secure the network connection using a TLS/SSL certificate. iDRAC's web server has a self-signed TLS/SSL certificate by default. The self-signed certificate can be replaced with a custom certificate, a custom signing certificate, or a certificate signed by a well-known Certificate Authority (CA). Whichever method you choose, when iDRAC is configured and the TLS/SSL certificate is installed on the management stations, TLS/SSL-enabled clients can access iDRAC securely and without certificate warnings.

SSH

iDRAC provides user control over the cryptographic settings for the SSH daemon such that you can determine the ideal settings for your environment. The control given to you is not a relaxation of the settings. Instead, the feature enables you to modify the value set for each option to achieve a narrower and stringent cryptographic policy. That is, you can only remove values from the options but cannot add any values other than those values that have been defined and allowed in the default value-set.

The cryptographic policies are configured using the following options:

- Ciphers Ciphers
- Host-Key-Algorithms HostKeyAlgorithms
- Key-Exchange Algorithms KeyExchangeAlgorithms
- MACs MACs

Typically, the values for each of these options are set to prudent settings that reflect the best security practices that cater to a wide variety of environments. As such, the iDRAC default settings for these options are the same as those options assigned by the SSH package open-source community. These settings can be configured using the RACADM command-line interface. See the *iDRAC RACADM CLI User's Guide*.



Automatic certificate renewal

iDRAC9 v4.0 and later has added a client for Simple Certificate Enrollment Protocol (SCEP) support and requires a Datacenter License. SCEP is a protocol standard used for managing certificates for large numbers of network devices using an automatic enrollment process. iDRAC can now integrate with SCEP-compatible servers such as the Microsoft ServerNDES service to maintain SSL/TLS Certificates automatically. This feature can be used to enroll and refresh a soon-to-be-expired web server certificate. You can use Server Configuration Profile to set the certificates on a one-to-one basis in the iDRAC UI. Also, you can provide scripts using tools such as RACADM.





iDRAC Cipher Select

The Cipher Suite Selection can be used to limit the ciphers that the web browser can use to communicate with iDRAC. Also, it can determine the security of the connection. These settings can be configured through the iDRAC web interface, RACADM, and Redfish. This functionality is available across several iDRAC releases – iDRAC7, iDRAC8 (2.60.60.60 and higher), and the current iDRAC9 (3.30.30.30 and higher).

Commercial National Security Algorithm (CNSA) support

The supported ciphers available in iDRAC9 with TLS1.3 Bit and 256 Bit Encryption are shown in the following figure. The ciphers available are inclusive of the ciphers in the CNSA-approved set.

Support	ed Server	Cipher(s)		
Preferred	TLSv1.2	256 bits	ECDHE-RSA-AES256-GCM-SHA384	Curve P-256 DHE 256
Accepted	TLSv1.2	256 bits	ECDHE-RSA-AES256-SHA384	Curve P-256 DHE 256
Accepted	TLSv1.2	256 bits	ECDHE-RSA-AES256-SHA	Curve P-256 DHE 256
Accepted	TLSv1.2	256 bits	DHE-RSA-AES256-GCM-SHA384	DHE 2048 bits
Accepted	TLSv1.2	256 bits	DHE-RSA-AES256-SHA256	DHE 2048 bits
Accepted	TL5v1.2	256 bits	DHE-RSA-AES256-SHA	DHE 2048 bits
Accepted	TLSv1.2	256 bits	DHE-RSA-CAMELLIA256-SHA	DHE 2048 bits

Tips for securing iDRAC connection

The most secure network connection is the iDRAC's Dedicated NIC because it can be connected to a network that is physically separated from the production network. This method physically separates the iDRAC management traffic from the production network traffic.

If use of the iDRAC's Dedicated NIC is not feasible, iDRAC can be run in Shared LOM mode with a VLAN enabled. However, the iDRAC's management traffic is sent across the same connection as the production network. Alternatively, if the use of a VLAN is not possible while in Shared LOM mode, access to iDRAC must be secured using strong passwords and other security measures.

For further information, see the iDRAC Security Configuration Guide.

IEEE.802.1x

The network standard IEEE 802.1x has been enabled on PowerEdge servers. This network protocol provides port-based network authentication. Devices requesting access to the LAN or WLAN must be authenticated and validated before authorization for access.

Domain isolation

The 14th, 15th, and 16th generations of PowerEdge servers provide additional security by using domain isolation, an important feature for multitenant hosting environments. To secure the server's hardware configuration, hosting providers might want to block any reconfiguration by tenants. Domain isolation is a configuration option that ensures that management applications in the host operating system have no access to the out-of-band iDRAC or to Intel chipset functions, such as Management Engine (ME) or Innovation Engine (IE).

Protecting data in use

Protecting application data that is being used in memory has become increasingly important. Whether data in use is a machine learning dataset or relates to keeping a secret in memory such as in multitenant environments, data in-use can be vulnerable to

threat vectors that can intrude on the contents of memory or the access bus. Data in use protection is necessary to secure computations that are increasingly operating on large datasets in memory. Additionally, code running on the data must be trusted and tamper-free. There must be facilities to separate trusted and nontrusted code execution environments for data in use.

New developments in CPU technology for confidential computing allow secure enclaves to protect application data at the hardware layer, enabling a more comprehensive data protection strategy. Starting with the 15th generation of server offerings, Dell Technologies enables these CPU technologies, including Intel SGX/TME and AMD SEV/SME.

AMD confidential compute features

AMD introduced Secure Encrypted Virtualization (SEV) with the first generation of AMD EPYC processors. It encrypts full system memory and individual virtual machine (VM) memory isolating the VM memory from the hypervisor. With each generation of AMD EPYC processors, AMD has enhanced SEV with additional features to safeguard privacy and integrity by encrypting each VM with one of up to 509 unique encryption keys known only to the processor. The keys are used to encrypt memory using 128-bit AES encryption engines in the memory controller. The hypervisor manages the keys in the memory controller with the help of AMD's Secure processor. The key AMD SEV technologies that are part of the Infinity Guard technology solution suite from AMD that address different use cases, deployments, and threat models include:

- AMD's SME technology refers to using a single key to encrypt system memory.
- AMD SEV uses one key per VM to isolate guests and the hypervisor from one another.
- AMD SEV-ES encrypts all CPU register contents when a VM stops running to prevent leakage of information from CPU registers to the hypervisor.
- AMD SEV-SNP adds strong memory integrity protection to help prevent hypervisor-based attacks like data replay, memory remapping, and so on.

Intel confidential computing features

Starting with the 3rd Generation Intel Xeon platform, Intel introduced several key security innovations. Total Memory Encryption (TME) is available to ensure memory accessed from the CPU is encrypted. By encrypting all memory, existing software applications run unmodified while simultaneously providing greater protection for system memory.

Intel TME helps ensure that all memory accessed from the Intel CPU is encrypted, including customer credentials, encryption keys, and other intellectual property or personal information about the external memory bus. Intel developed this feature to provide greater protection for system memory against hardware attacks, such as removing and reading the dual inline memory module (DIMM) after spraying it with liquid nitrogen or installing purpose-built attack hardware. Using the NIST storage encryption standard, AES XTS, an encryption key is generated using a hardened random number generator in the processor without exposure to software. This method allows existing software to run unmodified while better protecting memory.



Identity Access Management

Identity and Access Management (IAM) is a set of security controls to manage digital identities and provide authentication and authorization by controlling the requestor's access to information and resources, at the right level, and at the right time to limit unauthorized access. IAM frameworks provide advantages over simple passwords, such as:

- Enhanced security—Includes various tools such as single sign-on (SSO) services, multifactor authentication (MFA), and privileged access management for stronger security and to avoid risks from vulnerable passwords. IAM plays a critical role in protecting organizations against phishing and social engineering attacks by implementing strong authentication, access control, monitoring, and user education. It helps organizations reduce their attack surface and respond effectively to threats.
- **Reduction of compromised passwords**—Deters credential theft from phishing, social engineering, and brute-force attacks by using MFA and additional security layers.
- **Granular access control** Enables fine-grained control over user access permissions by using role-based access control (RBAC) and attribute-based access control (ABAC) to ensure that users only have access to necessary resources and data.
- **Centralized management**—Controls user accounts and access policies, which provides easier management of users as they join, move within, or leave the organization.
- Auditing and logging—Monitors user activities to identify and remediate suspicious access attempts.
- Scalability-Adapts to the need for a growing number of users and resources
- Regulatory compliance—Adheres to regulatory requirements to avoid legal consequences.
- User experience—Provides SSO services and password management to enable users to access multiple applications and services by using one set of credentials.
- Adaptive authentication—Applies additional security measures by assessing risk factors.
- Emergency access and recovery—Grants temporary emergency access for critical situations while maintain security controls.
- Integration—Integrates with various systems, applications, and cloud services for seamless access management across the entire IT ecosystem.

As a broad IT issue spanning technological and regulatory requirements, IAM is a strategic business imperative for all organizations to enhance security and resiliency. Zero Trust Architecture (ZTA) has emerged as the standard choice for securing all levels of infrastructure and is a foundational part of ZTA. The unofficial mantra of Zero Trust is "never trust, always verify" and if you cannot verify then you cannot trust. IAM provides that verification.



The foundation to a strong cybersecurity framework and the adoption and implementation of Zero Trust principles is identity. Identification of not only people but of applications, communication paths, network devices, and the data itself. If an organization does not have a strong Identity Credentialing and Access Management (ICAM) practice, then the underlying security practices are at risk. An effective identity and access solution must include necessary tools and controls that can capture and store user login details, facilitate the assignment and revocation of user access credentials, and oversee the central enterprise database of user roles, levels, and access privileges. At a minimum, access must only be allowed to facilitate the necessary functions that the organization has defined. As an example, users must only have access and permission to data, applications, and services to do what is defined in their job. Practices and tools must be implemented to ensure that this role is maintained and that even in that role, a user must be monitored to ensure that they are not doing anything that violates the organization's goals for data utilization.

MFA – Smartcards (CAC/PIV)

MFA for Smartcards (CAC/PIV) is a general-purpose certificate authentication that includes Common Access Card (CAC) and Personal Identification Verification (PIV) cards. Certificate authentication uses the client identity certificate to authenticate the user. It is used primarily in government or organizations that work with the federal industry.

MFA – RSASecurID

RSA SecurID is another means of authenticating a user on a system. As another twofactor authentication, iDRAC9 supports RSA SecurID with the Datacenter license and starting with firmware 4.40.00.00 and later, as another two-factor authentication.

Directory integration for authorization

For centralized user and domain management, iDRAC supports integration into privileged management tools such as Lightweight Directory Access Protocol (LDAP) and Active Directory. Using a directory service provides a central location for easier user inventory management and assigning user account access controls and settings.

You can use LDAP to authenticate users and groups in iDRAC. To configure the LDAP directory service, you can use the objects in the cfgLdap and cfgLdapRoleGroup groups with the config command. You can also use the objects in the iDRAC.LDAP and iDRAC.LDAPRole groups with the set command 12.

For Active Directory integration, you can configure LDAP over SSL (LDAPS) on iDRAC to communicate with Active Directory domain controllers. After a valid certificate is installed on the domain controller and the connection to the DC using SSL over port 636 is verified, you can use the directory service integration test on iDRAC/OME to communicate with the domain controller.

SSO

iDRAC supports SSO, which provides the ability to share validated credentials and identifications across multiple domains without having to rechallenge or reauthenticate the user. SSO enables an authenticated operating system administrator to directly access the iDRAC web interface without requiring login using separate iDRAC administrator credentials. iDRAC supports the following SSO protocols:

• **OpenID connect** is an open standard and is a decentralized authentication protocol that is typically used for machine-to-machine-like RestAPIs.



• Open-standards **Security Assertion Markup Language (SAML)** is typically used for UI SSO.

Role-based access controls

Role-based access control (RBAC) is the most popularly used form of access control. Permissions are grouped in roles and are typically assigned to a group. Assigning authorization capabilities to users and groups is managed in Active Directory.

You can set up user accounts with specific privileges (role-based authority to manage your system using iDRAC and maintain system security). By default, iDRAC is configured with a local administrator account. The default iDRAC username and password are provided with the system badge. As an administrator, you can set up user accounts to allow other users to access iDRAC. For more information, see the documentation for the server.

You can set up local users or use directory services such as Microsoft Active Directory or LDAP to set up user accounts. Using a directory service provides a central location for managing authorized user accounts.

iDRAC supports role-based access to users with a set of associated privileges. The roles are administrator, operator, read-only, or none. The role defines the maximum privileges available.

Time-based access controls

Time-based access control is another valuable tool for enhancing security and managing access to sensitive data, facilities, or systems in a tightly controlled and monitored environment. For instance, if a technician needs physical access to the USB port on the server, the iDRAC administrator can enable/disable specific access times for USB port access.

Scope-based access control

Scope-based access control provides more granular control than user- and role-based access control. It allows the administrator to apply a policy with a set of permissions that are evaluated when an entity tries to access the resource. For instance, access can be restricted, based on resource location, IP address range, and so on.

Capabilities and automation for efficient at-scale deployment

Zero-touch automation with iDRAC – Server configuration profiles

Zero Touch provisioning is available with iDRAC Enterprise or Datacenter licenses. Zero Touch provisioning automates all hardware configuration, certificate installation, repository firmware updates, and operating system deployment. The IT Admin can preconfigure security settings and ensure uniform server images. Zero Touch provisioning is available through the iDRAC Server Configuration Profile feature and with OpenManage Enterprise.

CloudIQ

CloudIQ monitors the health and cybersecurity of your enterprise-wide servers and predicts their performance so that you can proactively address issues before they impact your business.

CloudIQ offers a simple and intuitive solution to collect firmware details from PowerEdge servers including BIOS, iDRAC, NICs, PERC, drives, and supported peripherals. A recent

feature from both OpenManage Enterprise and CloudIQ identifies BIOS and firmware that requires an update. CloudIQ can report the current installed version, compare it to the latest Dell release available, and schedule updates. This information is collected from each server using the agent-free iDRAC, consolidated by OpenManage Enterprise, and then transferred to CloudIQ to be processed. This powerful feature includes user rights integrated in both OpenManage Enterprise and CloudIQ to allow only authorized users to run these commands. Also, CloudIQ can consolidate multiple OpenManage Enterprise instances into one server fleet management view.

Fourth Stage – Security management and monitoring

Challenges

You cannot defend what you cannot see. Attacks happen quickly, faster than a human can detect and respond. Protecting your critical assets in real-time within a dynamic and complex threat environment is a challenge. Lack of skills and resources exacerbates the problem. As IT administrators face ongoing challenges in their environments, more easily managing the infrastructure—through more automation, fewer task steps, and more intuitive interactions—is key to administration productivity.

It is critical that your business remains resilient and unaffected by adverse outcomes. The business must stay nimble to be aware, respond, and recover:

- Observability and transparency enable effective security event awareness and timely remediation.
- Lack of skills and resources increases problems. Monitor activity across infrastructure (user, device, data, network, and application) to identify patterns and anomalies. Analytics can be used to detect and respond to security threats.
- Automation is important.
- Analyze events, activities, and behaviors to derive context and apply Al/ML to achieve models that improve detection and reaction time in making real-time access decisions. Simplification is foundational and fundamental to a resilient infrastructure.

PowerEdge Solutions

The Dell management portfolio simplifies administrator tasks. It improves security and health monitoring to scale security confidently by using automation and intelligence. Monitoring and logging are a key part of a zero-trust implementation.

Visibility, logging, and alerts

It is critical to have a detection capability that provides complete visibility into the configuration, health status, and change events in a server system. This visibility must also detect malicious or other changes to BIOS, firmware, and option ROMs in the boot and operating system runtime process. Proactive polling must be coupled with the ability to send alerts for any events in the system. Logs must provide complete information about access and changes to the server. Most importantly, the server must extend these capabilities to all components.



Dell OpenManage Enterprise enables users to set up alert policies once and then automatically assign them for future alerts. Also, OpenManage Enterprise can apply a template to many servers at once. The OpenManage Enterprise solution ultimately saves time and effort by automating actions based on alerts after administrators have created a policy.

Telemetry

Beginning with iDRAC9 v4.00.00.00 firmware and a Datacenter license, IT managers can integrate advanced server hardware operation telemetry into their existing analytics solutions. Telemetry is provided as granular, timeseries data that is streamed or pushed. The advanced agent-free architecture in iDRAC9 provides over 180 data metrics that are related to server and peripheral operations. Metrics are precisely timestamped and internally buffered to allow highly efficient data stream collection and processing with minimal network loading. This comprehensive telemetry can be fed into analytics tools to predict failure events, optimize server operations, and enhance cyber resiliency. The iDRAC9 Telemetry Streaming collects and streams live system data from one or more PowerEdge servers to a centralized collector.

iDRAC Lifecycle Logs

The Lifecycle Logs are a collection of events that occur in a server over time. They provide a description of events with timestamps, severity, user ID or source, and recommended actions. This technical information helps with security tracking and other hardware alerts.

The various types of information that is recorded in the Lifecycle Controller Log (LCL) include:

- Configuration changes on the system hardware components
- iDRAC, BIOS, NIC, and RAID configuration changes
- Logs of all the remote operations
- Firmware update history based on device, version, and date
- Information about replaced parts
- Information about failed parts
- Event and error message IDs
- Host power-related events
- POST errors
- User login events
- Sensor state change events

Alerts

iDRAC provides the capability to configure different event alerts and actions to be performed when a Lifecycle Logs event occurs. When an event is generated, it is forwarded to the configured destinations by using the selected alert type mechanisms. Users can enable or disable alerts through the iDRAC web interface, RACADM, or with the iDRAC settings utility.



iDRAC supports several types of alerts such as:

- Email or IPMI alert
- SNMP trap
- Operating system and Remote System logs
- Redfish event

Alerts are categorized by severity – Critical, Warning, or Informational. The following filters can be applied to alerts:

- System health-For example, temperature, voltage, or device errors
- Storage health-For example, controller errors, physical or virtual disk errors
- **Configuration changes**—For example, change in RAID configuration, PCIe card removal
- Audit logs-For example, password authentication failure
- **Firmware**—For example, upgrades or downgrade

The IT administrator can set different actions for alerts – Reboot, Power Cycle, Power Off, or No action.

TLS for Remote Syslog

The iDRAC Remote Syslog feature allows you to write the RAC log and the System Event Log (SEL) remotely to an external syslog server. You can read all logs from the entire server farm from a central log. The Remote Syslog protocol does not require user authentication. For the logs to be entered in the Remote Syslog server, ensure that there is proper network connectivity between iDRAC and the Remote Syslog server and that the Remote Syslog server is running on the same network as iDRAC.

The iDRAC's web server has a self-signed TLS/SSL certificate by default. The selfsigned certificate can be replaced with a custom certificate, a custom signing certificate, or a certificate signed by a well-known Certificate Authority (CA). Redfish scripts can perform automated certificate uploads. When a link has been established between the two servers, TLS encryption and SSL decryption enable secure data transport.

SELinux framework

SELinux operates at the core kernel level on iDRAC and does not require user input or configuration. SELinux logs security messages when an attack is detected. These log messages indicate when and how an attacker tried to break into the system. These logs are available through SupportAssist to customers enrolled in this new feature. In future releases of iDRAC, these logs will be available in the Lifecycle Controller Logs.

Real-time detection – BIOS Live scan

BIOS live scanning verifies the integrity and authenticity of the BIOS image in the primary ROM when the host is powered on. BIOS live scanning is not in the POST process. This feature is available only with the iDRAC9 4.10.10.10 (supported AMD platforms) and iDRAC9 4.40.20.00 (supported Intel platforms) Datacenter licenses. You must have administrator privileges or operator privileges with the "Execute Debug Commands" debug privilege to perform this operation.



You can initiate BIOS image scanning either on demand or by scheduling the scan through the iDRAC UI, RACADM, and Redfish interfaces. The BIOS live scan feature is available starting with 15th generation of PowerEdge servers with AMD "Rome"-based processors or Intel "Ice Lake"-based processors.

Boot time and run time BIOS scanning

A critical aspect of server security is ensuring that the boot process is verified as secure. This process provides a trusted anchor for all subsequent operations such as booting an operating system or updating firmware. PowerEdge servers have used silicon-based security for several generations for features such as iDRAC Credential Vault, an encrypted secure memory in iDRAC for storing sensitive data. The boot process is verified using a silicon-based RoT to meet the following recommendations:

- NIST SP 800-147B, BIOS Protection Guidelines for Servers
- NIST SP 800-155, BIOS Integrity Measurement Guidelines

On the Dell PowerEdge servers with iDRAC9, iDRAC first boots with chain of trust authentication, and then verifies BIOS integrity. iDRAC assumes the role of hardwarebased root of trust. For AMD platforms, iDRAC accesses the primary BIOS ROM through SPI and the AMD fusion controller hub (FCH), and performs the RoT process. For Intel platforms, iDRAC accesses the primary BIOS ROM through SPI and the Intel Platform Controller Hub (PCH), and performs the RoT process.

iDRAC9 directly accesses the BIOS primary ROM to perform a RoT operation on the processor on both the security block and the host Initial Boot Block.

Silicon-based Root of Trust

PowerEdge servers use an immutable, silicon-based RoT to attest to the integrity of BIOS and iDRAC9 firmware cryptographically. This RoT is based on one-time programmable, read-only public keys that provide protection against malware tampering. The BIOS boot process uses Intel Boot Guard technology or AMD Platform Secure Boot technology. This technology verifies that the digital signature of the cryptographic hash of the boot image matches the signature stored in silicon by Dell Technologies in the factory. A verification failure results in a server shutdown and user notification in the Lifecycle Controller Log. The user can initiate the BIOS recovery process. If Boot Guard validates successfully, the other BIOS modules are validated by using a chain of trust procedure. Then, control is given to the operating system or hypervisor. In addition to Boot Guard, iDRAC9 4.10.10.10 or later provides a RoT mechanism that verifies the BIOS image at the host boot time. The host can boot only after the BIOS image is successfully validated. iDRAC9 also provides a mechanism to validate the BIOS image at run time, on demand, or at user-scheduled intervals.

For the chain of trust, each BIOS module contains a hash of the next module in the chain. The key modules in the BIOS include:

- Technical support and resources ID 483
- Initial Boot Block (IBB)
- Security (SEC)
- Pre-EFI Initialization (PEI)
- Memory Reference Code (MRC) o Driver Execution Environment (DXE)



• Boot Device Selection (BDS)

If Intel Boot Guard authenticates the IBB module, the IBB module validates the SEC and PEI modules before handing control to it. The SEC and PEI modules then validate the PEI and MRC modules, which further validates the DXE and BDS modules. Next, control is handed over to UEFI Secure Boot. Similarly, for PowerEdge AMD EPYC-based servers, AMD Secure Root of Trust technology ensures that servers boot only from trusted firmware images. AMD Secure Run Technology encrypts the main memory, keeping it private from malicious intruders accessing the hardware. There are no required application modifications to use this feature, and the security processor never exposes the encryption keys outside of the processor. iDRAC takes on the role of hardware-based security technology and accesses the primary BIOS ROM through SPI. iDRAC, along with the AMD fusion controller hub (FCH) performs the RoT process.

Under the following conditions, iDRAC9 recovers the BIOS:

- BIOS integrity check failed
- BIOS self-check failed

Note: Use the RACADM command to recover the BIOS setup.

The iDRAC boot process uses its own independent silicon-based RoT that verifies the iDRAC firmware image. The iDRAC RoT also provides a critical trust anchor for authenticating the signatures of Dell firmware update packages (DUPs).

System lockdown

iDRAC9 offers a feature that 'locks down' the server hardware and firmware configuration and requires an Enterprise or Datacenter license. You can enable this mode by using the UI, the RACADM CLI, or the Server Configuration Profile. Users with administrative privileges can set System Lockdown mode, which prevents users with lesser privileges from changing the server. The IT administrator can enable or disable this feature. Any changes made when System Lockdown is disabled are tracked in the Lifecycle Controller Log. By enabling lockdown mode, you can prevent configuration drift in your data center when using Dell tools and agents and protect against malicious attacks against embedded firmware when using Dell Update Packages. Lockdown mode can be enabled dynamically, without requiring a system reboot. iDRAC9 v4.40 introduced enhancements where in addition to the current System Lockdown which only controls the updates using Dell Update Package (DUP), this lockdown functionality is extended to select NICs.

Note: Enhanced Lockdown for NICs only includes firmware lockdown to prevent firmware updates.

Configuration (x-UEFI) lockdown is not supported. When the customer sets the system in lockdown mode by enabling or setting attributes from any of the supported interfaces, iDRAC will take additional actions depending on the system configuration. These actions depend on the third-party devices detected as part of the iDRAC discovery process.

Drift detection

By enforcing standardized configurations and adopting a "zero tolerance" policy for any changes, organizations can reduce the potential for exploitation. Dell OpenManage



Enterprise Console allows you to define your own server configuration baselines and then monitor the drift of your production servers from those baselines. The baseline can be built based on different criteria to fit different production enforcement, such as security and performance.

OpenManage Enterprise can report any deviations from the baseline and optionally repair the drift with a simple workflow to stage the changes on iDRAC out of band. The changes can then take place at the next maintenance window while servers reboot to make the production environment compliant again. This staged process enables you to deploy configuration changes to production without any server downtime during nonmaintenance hours. It increases the server availability without compromising the serviceability or security.

Chassis intrusion detection

PowerEdge servers provide hardware intrusion detection and logging, with detection working even when no AC power is available. Sensors on the chassis detect when anyone opens or tampers with the chassis, even during transit. Servers that have been opened while in transit generate an entry in the iDRAC Lifecycle Logs after power is supplied.

Automated and manual recovery

BIOS and operating system recovery

The 14th, 15th, and 16th generations of PowerEdge servers include two types of recovery: BIOS Recovery and Rapid Operating System (OS) Recovery. These features enable rapid recovery from corrupted BIOS or operating system images. In both cases, a special storage area is hidden from run-time software (BIOS, operating system, device firmware, and so on). These storage areas contain pristine images that can be used as alternatives to the compromised primary software.

Rapid Operating System (OS) Recovery enables rapid recovery from a corrupted operating system image (or an operating system image suspected of malicious tampering). The recovery media can be accessed using an internal SD card, SATA ports, M.2 drives, or internal USB. The selected device can be exposed to the boot list and the operating system to install the recovery image. It can then be disabled and hidden from the boot list and operating system. In the hidden state, the BIOS disables the device so that the operating system cannot access it. If there is a corrupted operating system image, the recovery location can then be enabled for the boot process. These settings can be accessed through BIOS or the iDRAC interface.

In extreme cases, if the BIOS is corrupted (either by a malicious attack, a power loss during the update process, or any other unforeseen event), it is important to provide a way to recover the BIOS to its original state. A backup BIOS image is stored in iDRAC so it can be used to recover the BIOS image if needed. iDRAC orchestrates the entire end-to-end recovery process.

- BIOS itself initiates automatic BIOS recovery.
- Users using the RACADM CLI command can initiate on-demand BIOS recovery.

FW rollback

We recommend that you update the firmware to ensure you have the latest features and security updates. However, you may need to rollback an update or install an earlier

version if you encounter issues after an update. If you roll back to the previous version, it is also verified against its signature.

Firmware Rollback from existing production version "N" to a previous version "N-1" is supported for the following firmware images:

- BIOS
- iDRAC with Lifecycle Controller
- Network Interface Card (NIC)
- PowerEdge RAID Controller (PERC)
- Power Supply Unit (PSU)
- Backplane

You can roll back the firmware to the previously installed version (N-1) using any of the following methods:

- iDRAC web interface
- CMC web interface
- RACADM CLI for iDRAC and CMC
- Lifecycle Controller UI
- Lifecycle Controller remote services

You can roll back the firmware for iDRAC or any device that Lifecycle Controller supports, even if the upgrade was previously performed using another interface. For example, if the firmware was upgraded using the Lifecycle Controller UI, you can roll back the firmware using the iDRAC web interface. You can perform firmware rollback for multiple devices with one system reboot.

On 14th, 15th, and 16th generations of PowerEdge servers that have a single iDRAC and Lifecycle Controller firmware, rolling back the iDRAC firmware also rolls back the Lifecycle Controller firmware.

Firmware rollback protection

If the firmware has a known vulnerability which would expose your server to attack, the BIOS itself can prevent downgrade to a previous version. The firmware release notes state that you cannot rollback when performing an update.

Full Power Cycle

In a Full Power Cycle, the server and all its components are rebooted. It drains main and auxiliary power from the server and all components. All data in volatile memory is also erased.

A physical Full Power Cycle requires removing the AC power cable, waiting for 30 seconds, and then putting the cable back. This method poses a challenge when working with a remote system. A feature in the 14th, 15th, and 16th generations of PowerEdge servers enables you to perform an effective full power cycle from the iDRAC Service Module (iSM), IDRAC UI, BIOS, or a script. A full power cycle takes effect at the next power cycle.



The Full Power Cycle feature eliminates the need for anyone to be physically present in the data center, thus reducing time to troubleshoot. It can eliminate, for example, any malware that is still memory-resident.

Updating

Dell Technologies provides a rich set of tools to make it easier to keep your server's firmware and BIOS up-to-date and update quickly. Ensuring firmware is up to date is a vital task to keep production servers secure and operating efficiently. Tracking and implementing these updates can be burdensome for administrators. iDRAC9 provides automatic updates with the ability to schedule firmware updates as wanted.

Many companies schedule monthly maintenance windows to handle operating system, application, and firmware updates. With OpenManage Enterprise, systems administrators can stage firmware updates for the next time the system is rebooted or for a scheduled deployment. This method ensures that nobody must be physically present to run the updates.

PowerEdge provides patches to security vulnerabilities with Dell security advisories. The advisories provide timely information, guidance, and mitigation options to minimize risks associated with security vulnerabilities.

Restoring server configuration after hardware servicing

Remediating service events is a critical part of any IT operation. The ability to meet recovery time objectives and recovery point objectives has direct implications on the security of the solution. Restoring server configuration and firmware assures that security policies for server operation are automatically met.

PowerEdge servers provide functionality that quickly restores server configuration in the following situations:

- Individual part replacement
- Motherboard replacement (full server profile backup and restore)
- Motherboard replacement (easy restore)

Parts replacement

iDRAC automatically saves the firmware image and configuration settings for NIC cards, RAID controllers, and Power Supply Units (PSUs). If there is a field replacement of these parts, iDRAC automatically detects the new card and restores the firmware and configuration to the replaced card. This functionality saves critical time and ensures a consistent configuration and security policy. The update occurs automatically on system reboot after replacing the supported part.

Easy Restore (for motherboard replacement)

Easy Restore is an integrated storage component that maintains critical configuration information. The Easy Restore feature allows you to restore your system's service tag, all licenses, UEFI configuration, system configuration settings (BIOS, iDRAC and NIC), and the OEM ID (Personality Module) after replacing the system board. All data is backed up to a backup flash device automatically. If the BIOS detects a new system board and the service tag in the backup flash device, the BIOS prompts you to restore the backup information. You can still choose to perform the full system backup with iDRAC8 as you do with iDRAC6 and iDRAC7. This solution backs up and restores the actual firmware

versions in addition to the hardware settings. Easy Restore does not copy the firmware drivers due to size limitations.

CloudIQ

Misconfigurations of infrastructure systems can open your organization to cyber intrusion and are a leading threat to data security. The CloudIQ cybersecurity feature proactively monitors infrastructure security configurations for Dell PowerStore and PowerMax storage systems and PowerEdge servers, and notifies users of security risks. A risk level is assigned to each system, placing the system into one of four categories, depending on the number and severity of the issues: Normal, Low, Medium, or High.

By using CloudIQ Cybersecurity policy templates, users can quickly set up security configuration evaluation tests and assign them to large numbers of systems with just a few clicks. Once assigned, the test plan is evaluated against each associated system, and the system administrator is notified in minutes of any unwanted configuration settings.

When a security risk is found, remediation instructions are provided to help you address the issue quickly. CloudIQ evaluates outgoing Dell Security Advisories (DSAs) and intelligently notifies users when those advisories are applicable to their specific Dell system models with specific system software and firmware versions. This notification eliminates the need for users to investigate if a Security Advisory applies to their systems and allows them to focus on remediation immediately.

Managed detection and response services

Dell Technologies Managed Detection and Response services is a cloud-based offering that helps organizations quickly and significantly improve their security posture—while reducing the burden on IT. This fully managed, end-to-end, 24/7 service monitors, detects, investigates, and responds to threats across the entire IT environment.

Designed for organizations with 50 endpoints or more, this unique service uses two key capabilities:

- The power of the open Secureworks Taegis XDR security analytics software, built on more than 20 years of SecOps expertise including real-world threat intelligence and research, and experience detecting and responding to advanced threats
- The expertise of Dell Technologies security analysts, gained through years of experience helping organizations worldwide to better protect their business

<u>Dell Technologies Managed Detection and Response</u> provides around-the-clock access to security experts. They provide end-to-end visibility and protection from endpoint to cloud, covering every aspect of advanced threat detection and response supported by Taegis XDR's database of 52,000 unique threat indicators that are managed and updated daily. Taegis XDR also ingests data from existing security solutions to use any previous security investments.

Dell Technologies security analysts assist with initial setup, monitoring, detection, remediation, and response—all for one predictable price. They work closely with your IT team to understand the environment. They provide advice about improvements to the security posture and help you set up and deploy the Taegis XDR software agent to endpoints. Then, using the Taegis XDR application, they monitor and review alerts 24/7.



Fifth stage - Secure decommissioning and repurposing

If an alert merits investigation, analysts determine and perform the appropriate response. If a threat is malicious or requires your action, you are informed and, if necessary, provided with step-by-step instructions. As part of the service, Dell Technologies also provides up to 40 hours per quarter of <u>remote remediation assistance</u>, such as helping with troubleshooting, issue resolution, software deployments, patch and asset assessment, and configuration of IT environments.

If there is a security incident, Dell Technologies initiates the process to get your business up and running and provides up to 40 hours of remote incident response assistance a year.

Fifth stage – Secure decommissioning and repurposing

Challenge Data security is a key consideration throughout the life cycle of a server, including when the server is repurposed or retired. Many servers are repurposed as they are transitioned from workload to workload, or as they change ownership from one organization to another. All servers are retired when they reach the end of their useful life. When such transitions occur, IT best practices recommend removing all data from the server to ensure that confidential information is not inadvertently shared.

PowerEdge solutions

Beyond best practices, in many cases government regulations about privacy rights also necessitate complete data elimination when IT resources are transitioned. Data erasure is a key capability encompassed in the Dell Secure Development Lifecycle (SDL). The SDL and secure server management tools ensure that PowerEdge servers are secure at every stage in the server life cycle, from server conception, design and manufacturing, to operation and decommissioning.

At this final stage (decommissioning/retirement), or when a server is repurposed due to change of workload or ownership, a capability starting with the PowerEdge 14th generation of servers, can simplify data erase.

System Erase, with iDRAC9 and the 14th, 15th, and 16th generations of PowerEdge servers, simplifies the process of erasing server storage devices and server nonvolatile stores such as caches and logs. To meet varying Systems Administrator needs for interactive and programmable operations, the following methods can perform System Erase:

- Lifecycle Controller UI
- WS-Man API
- RACADM CLI

Using one of these methods, an administrator can selectively reset a PowerEdge server to its original state (factory settings), removing data from internal server nonvolatile stores and from storage devices within the server.

System Erase can discover server-attached storage including hard disk drives (HDDs), SEDs, ISE, and nonvolatile memory drives (NVMes). Data stored on ISE, SED, and NVMe devices can be made inaccessible using cryptographic erase while devices such as non-ISE SATA HDDs can be erased using data overwrite.

40 Cyber Resilient Security in Dell PowerEdge Servers

Inserido ao protocolo 22.951.206-4 por: Pedro Henrique Golin Linhares em: 31/03/2025 12:49. A autenticidade deste documento pode ser validada no endereço: https://www.eprotocolo.pr.gov.br/spiweb/validarDocumento com o código: fc1c44dfe0686d5547ffd691199be699.



Secure Erase

Through the life cycle controller, customers can repurpose or retire a system. All drives now shipping on PowerEdge systems can be securely erased. On an older platform that might not have encryption-capable drives, there is a "standard disks (overwrite data)" option. Unrecoverable processes generate warning messages. The server is powered off after a retire or repurpose operation. You can view the Lifecycle Logs in iDRAC to confirm that the operation was successful.

At the end of a system life cycle, it can either be retired or repurposed. For either scenario, System Erase removes sensitive data and settings from the server. Secure Erase wipes storage devices and server nonvolatile stores such as caches and logs so that no confidential information unintentionally leaks. It is a utility in Lifecycle Controller (F10) that erases logs, configuration data, storage data, and cache.

The System Erase feature can erase the following devices, configuration settings, and applications:

- iDRAC is reset to default settings, erasing all data and settings.
- Lifecycle Controller (F10) data is cleared.
- BIOS and NVRAM are reset to default settings.
- Embedded diagnostics and operating system driver packs are cleared.
- iDRAC Service Module (iSM) are cleared.
- SupportAssist Collection reports are cleared.

The following components can also be erased:

- Hardware Cache (clear PERC NVCache)
- vFlash SD Card (initialize card)

Note: vFlash not available on 15th generation of PowerEdge servers or later.

System Erase cryptographically disposes of data on the following components:

- Self-Encrypting Drives (SED)
- ISE drives
- NVM devices such as Intel Apache Pass and NVDIMMs

Secure erase – physical disk

Reset the drive to the factory settings. All data on the SSD is permanently removed and cannot be recovered. To sanitize the drive, the mapping table is deleted and all blocks to which data has been written are erased. Not all SSDs support the sanitize feature.

Overwrite erase

Overwrite-erase is a software-based method that overwrites data with zeros and ones. Data overwrite can erase non-ISE SATA hard drives.



Summary

Crypto erase

ISE destroys the internal encryption key that is used in the 14th, 15th, and 16th generations of PowerEdge drives and renders the user data unrecoverable. Used on self-encrypting drives, the encryption key is erased. The data remains on the drive but is inaccessible without the key. ISE is a recognized method of data erasure on storage drives as seen in NIST Special Publication 800-88 "Guidelines for Media Sanitization."

Advantages of the new ISE feature with System Erase include:

- **Speed**—ISE is faster than data overwriting techniques like DoD 5220.22-M (seconds compare with hours)
- **Effectiveness**—ISE renders all the data on the drive, including reserved blocks, unreadable.
- Better TCO—Storage devices can be reused instead of being physically destroyed.

The following methods can perform System Erase procedure:

- Lifecycle Controller interface (F10)
- RACADM CLI
- Redfish

Data sanitization and destruction services

Data continues to grow and drive strategic advantage. Meanwhile, national security issues and data privacy regulations are escalating. As organizations navigate technology changes, they are further challenged with data security and compliance.

The end of the product life cycle is an aspect of data security that is increasingly important. Data Sanitization for Enterprise is a software-based method of securely overwriting data to render it unrecoverable. The various options include:

- Data Sanitization for Enterprise Onsite—An option for customers looking to refresh or redeploy assets. This service performs sanitization at the business' location, securing data while assets remain in the environment.
- Data Sanitization for Enterprise Offsite with Asset Resale and Recycle—A service that removes assets from the business' environment, sanitizes them at a secure location, and evaluates them for resale or reuse. The customer is compensated if value is found. If no value is found, assets are recycled according to local regulatory guidelines.

Data Destruction for Enterprise—A process that renders data inaccessible through the process of physical shredding. It is available for all Dell infrastructure solutions and similar third-party non-Dell branded assets. This process does not require systems to be operational.

Summary

Data center security is paramount to business success, and the security of the underlying server infrastructure is critical. Cyberattacks have the potential for extended system and business downtime, lost revenue and customers, legal damages and tarnished corporate

reputation. To protect, detect, and recover from hardware-targeted cyberattacks, security must be built into server hardware design, not added on later.

Dell Technologies has been a leader in using silicon-based security to secure firmware and protect sensitive user data in PowerEdge servers for the past two generations. The 14th, 15th, and 16th generations of PowerEdge product lines feature an enhanced Cyber Resilient Architecture that uses silicon-based Root-of-Trust to further harden server security including the following features:

- Cryptographically verified Trusted Booting that anchors end-to-end server and overall data center security. It includes features like silicon-based Root-of-Trust, digitally signed firmware, and automatic BIOS recovery.
- Secure Boot, which checks the cryptographic signatures of UEFI drivers and other code loaded before the operating system runs.
- iDRAC Credential Vault, a secure storage space for credentials, certificates, and other sensitive data that is encrypted with a silicon-based key that is unique for every server.
- Dynamic System Lockdown, a capability unique to PowerEdge servers, helps secure any system configuration and firmware from malicious or unintended changes while alerting users to any attempted system changes.
- Enterprise Key Management delivers a central key management solution to manage data-at-rest across the organization.
- System Erase, which allows users to easily retire or repurpose their 14th, 15th, and 16th generations of PowerEdge servers by securely and quickly wiping data from storage drives and other embedded nonvolatile memory.
- Supply Chain Security provides supply chain assurance by ensuring there is no product tampering or counterfeit components before shipping products to the customers.

The 14th, 15th, and 16th generations of PowerEdge servers, with their industry-leading security, form a trusted foundation for IT transformation on which customers can securely run their IT operations and workloads, and accelerate to Zero Trust adoption. Dell Technologies stops at nothing to help our customers build their breakthrough deployments. Our modern security approach ensures that an organization's environment is secure and resilient so that customers can focus on their core competencies, introduce their innovations, and advance human progress.



References

Dell Technologies documentation The following Dell Technologies documentation provides other information related to this document. Access to these documents depends on your login credentials. If you do not have access to a document, contact your Dell Technologies representative.

- iDRAC9 Security Configuration Guide
- Dell EMC Secured Component Verification Reference Guide for Servers
- <u>Understanding Confidential Computing with Trusted Execution Environments</u> and Trusted Computing Base models
- iDRAC9 System Lockdown: Preventing Unintended Server Changes
- Next Generation Dell PowerEdge Servers: Transition to Modern UEFI
- Dell EMC PowerEdge UEFI Secure Boot Customization: Reduce Attack Surface with Complete Control of Certificates
- Dell Technologies Supply Chain Security: Secured Component Verification for <u>PowerEdge</u>
- <u>Dell PowerEdge: iDRAC Automatic Certificate Enrollment</u>
- Improved iDRAC9 Security using TLS 1.3 over HTTPS on Dell PowerEdge Servers
- <u>A Partnership of Trust: Dell Supply Chain Security</u>
- <u>PowerEdge Advantages in your Zero Trust Journey Video</u>
- <u>AMD on PE Extending Data Protection to Data in Use Video</u>
- <u>AMD on PE Extended Boot Protection Video</u>
- <u>Zero Trust Architecture Video</u>
- Cyber Resilient Architecture Video
- <u>Secured Component Verification Video</u>
- <u>SEKM Video</u>
- IPv6 Direct from Development
- <u>iDRAC9 Virtual Power Cycle: Remotely power cycle Dell EMC PowerEdge</u> <u>Servers</u> -Direct from Development
- Transform Datacenter Analytics with iDRAC9 Telemetry Streaming
- <u>Configure iDRAC to use Active Directory Authentication (dell.com)</u>
- <u>Securing 14th Generation Dell EMC PowerEdge Servers with System Erase</u>
- (Direct from Development) Security in Server Design
- (Direct from Development) Cyber-Resiliency Starts At The Chipset And Bios
- Factory Generated Default Password for iDRAC9 for Dell EMC 14th Generation (14G) PowerEdge Servers

References

61000 Fls. <u>989</u> Mov. <u>33</u>

- Dell EMC iDRAC Response to Common Vulnerabilities and Exposures (CVE) CVE-2017- 1000251 "BlueBorne"
- <u>(Video) Secure Boot Configuration And Certificatemanagement Using</u> <u>RACADM- Video</u>
- Secure Boot Management on Dell EMC PowerEdge Servers
- <u>Signing UEFI images for Secure Boot feature in the 14th and 15th generation</u> and later Dell EMC PowerEdge servers
- <u>Rapid Operating System Recovery</u>
- <u>Managing iDRAC9 Event Alerts on 14th generation (14G) Dell EMC PowerEdge</u>
 <u>Servers</u>
- UEFI Secure Boot Customization
- iDRAC Overview
- OpenManage Console Overview
- OpenManage Mobile Overview
- <u>Motherboard Replacement</u>



ProSupport Plus for Infrastructure

Introdução

A Dell Technologies¹ tem o prazer de fornecer o ProSupport Plus for Infrastructure (os "Serviços" ou "Serviços de suporte") em conformidade com esta Descrição de serviço ("Descrição de serviço"). A cotação, o formulário de pedido, outro formulário de fatura mutuamente acordado entre as partes ou a confirmação do pedido pela Dell Technologies (o "Formulário de pedido") incluirá o(s) nome(s) do(s) Produto(s)², o(s) Serviço(s) aplicável(is) e a(s) opção(ões) relacionada(s), caso existam. Para obter assistência adicional ou solicitar uma cópia do contrato vigente aplicável aos Serviços (o "Contrato"), entre em contato com um representante de vendas da Dell Technologies. Para obter uma cópia de seu contrato com o revendedor Dell Technologies aplicável, entre em contato com o revendedor.

O escopo deste Serviço

Os recursos deste Serviço incluem:

- Acesso por telefone 24 (vinte e quatro) horas por dia, 7 (sete) dias por semana (inclusive feriados)³ à central de especialistas globais da Dell Technologies, que é composta por engenheiros de nível sênior do ProSupport para assistência à solução de problemas de hardware e software.
- Envio no local de um técnico de campo especializado da Dell Technologies e/ou entrega de peças de substituição no Local de instalação ou em outra localização comercial indicada pelo Cliente e aprovada pela Dell Technologies, como detalhado no Contrato para resolver um problema no Produto.
- Acesso a um Service Account Manager (SAM).
- Substituição de peças, quando considerado necessário para resolver ou evitar um problema
- As plataformas AIOps do ProSupport incluem o APEX AIOps Infrastructure Availability, o TechDirect e o MyService 360, que são todos ativados pelo software de conectividade, como o Gateway de conexão segura (SCG), e oferecem benefícios não limitados a estes:
 - Detecção proativa de problemas e criação automatizada de casos
 - o Detecção preditiva de falhas de hardware
 - Criação de casos de autoatendimento
 - o Despacho de peças de autoatendimento
 - Avisos de segurança da Dell
 - Avaliação de segurança cibernética do APEX AIOps Infrastructure Availability

Consulte a tabela abaixo para obter mais informações.

Como entrar em contato com a Dell Technologies para solicitar um Serviço

Suporte on-line, por chat e por e-mail: O suporte da Dell Technologies por site, chat e e-mail está disponível para alguns produtos em <u>www.dell.com/contactus</u>.

Solicitação de suporte por telefone: Disponível 24x7 (inclusive em feriados). A disponibilidade pode variar fora dos Estados Unidos e limita-se a esforços comercialmente razoáveis, exceto quando especificado de outra forma neste documento. Acesse <u>www.dell.com/contactus</u> para obter uma lista de números de telefone aplicáveis à sua localização.

A tabela a seguir lista os recursos de serviço do ProSupport Plus for Infrastructure fornecidos de acordo com os termos da garantia e/ou manutenção da Dell Technologies. O ProSupport Plus for Infrastructure está disponível para dar suporte e manter:

1. Equipamentos da Dell Technologies identificados na <u>Tabela de garantia e manutenção de produtos da Dell Technologies</u> e/ou no Formulário de pedido como:

- Incluindo o ProSupport Plus for Infrastructure pela duração da garantia aplicável; ou

Dell Technologies ProSupport Plus for Infrastructure | v11 | May 2024

¹ "Dell Technologies", conforme usado neste documento, significa a entidade de vendas da Dell ("Dell") aplicável, especificada no Formulário de pedido da Dell, e a entidade de vendas da EMC ("EMC") aplicável, especificada no Formulário de pedido da EMC. O uso de "Dell Technologies" neste documento não indica alteração no nome legal da entidade Dell ou EMC com a qual você fez negócios.

² Conforme usado neste documento, "Produtos da Dell Technologies", "Produtos", "Equipamento" e "Software" significam o Equipamento e o Software da Dell Technologies identificados na <u>Tabela de garantia e manutenção de produtos da Dell Technologies</u> ou no seu Formulário de pedido, e "Produtos de terceiros" é definido no seu Contrato ou, na ausência de tal definição em seu Contrato, nos <u>Termos de venda comerciais da Dell Technologies</u> ou nos termos de venda locais da Dell Technologies, conforme aplicável. "Você" e "Cliente" referem-se à entidade nomeada como o comprador desses Serviços indicados no Contrato.

³ A disponibilidade varia de acordo com o país. Para obter mais informações, entre em contato com seu representante de vendas.


- Elegíveis ao upgrade para o ProSupport Plus for Infrastructure pela duração da garantia aplicável; ou
- Elegíveis ao ProSupport Plus for Infrastructure durante um período de manutenção subsequente.
- 2. Software da Dell Technologies identificado na <u>Tabela de garantia e manutenção de produtos da Dell Technologies</u> e/ou no Formulário de pedido como elegível ao ProSupport Plus for Infrastructure durante um período de manutenção.

RECURSO DO SERVIÇO	DESCRIÇÃO	PROSUPPORT PLUS — DETALHES DA COBERTURA		
SUPORTE TÉCNICO GLOBAL	O Cliente entra em contato com a Dell Technologies por telefone ou pela interface Web 24x7 para informar um problema no Equipamento ou no Software. Os contatos por telefone serão encaminhados para um contato de suporte técnico remoto para resolução do problema. A criação automatizada de casos é disponibilizada quando a AlOps platform é configurada por meio do Gateway de conexão segura (SCG).	 Incluso. Para problemas de Severidade 1, os clientes obtêm gerenciamento de escala e procedimentos "CritSit" para situações críticas com a cobertura do Gerente de incidentes. 		
RESPOSTA NO LOCAL	A Dell Technologies envia profissionais autorizados ao Local de instalação para que trabalhem no problema após a Dell Technologies ter isolado o problema e confirmado a necessidade da Resposta no local.	 Incluída apenas para Equipamento. O objetivo inicial da Resposta no local é uma resposta de serviço em quatro horas depois que a Dell Technologies considerar necessário a Resposta no local. <u>Resposta no local</u> O técnico geralmente chega ao local em 4 horas após a conclusão da solução de problemas e o isolamento do problema. Disponível 24 (vinte e quatro) horas por dia, 7 (sete) dias por semana, inclusive feriados. Disponível nos locais definidos com resposta em 4 (quatro) horas. O estoque de componentes operacionais essenciais está a no máximo 4 horas de distância da localização do cliente, conforme determinado pela Dell Technologies. As peças não essenciais podem ser enviadas por meio de entrega noturna. A resposta no local não será aplicável ao software e poderá ser 		
SUPORTE ESSENCIAL	Para problemas de Severidade 1, a Dell Technologies realiza a cobertura incluída, conforme considerado necessário pela Dell Technologies.	 Incluída apenas para Equipamento. Procedimentos de situação crítica — os problemas de Severidade 1 se qualificam à atuação rápida do Gerente de escala/resolução e à cobertura de incidentes "CritSit". Envio rápido: Despacho de engenheiro de campo sênior em conjunto com a solução de problemas por telefone. O engenheiro e a disponibilidade são determinados pela Dell conforme aplicável. Diagnóstico no local sob demanda da Dell quando a equipe do cliente não estiver disponível ou razoavelmente capaz de executar a solução de problemas no local. Somente para equipamento conectado ao gateway de conexão segura. A solicitação do cliente deve ser iniciada por meio de um chamado por telefone. Prioridade de produção em caso de situação crítica causada por um desastre natural. Em muitos casos, isso inclui a produção acelerada de um novo sistema da Dell Technologies. 		



RECURSO DO SERVIÇO	DESCRIÇÃO	PROSUPPORT PLUS — DETALHES DA COBERTURA
CHAMADA DE 6 HORAS PARA REPARO 24X7: RESPOSTA NO LOCAL EM 4 HORAS COM SERVIÇO DE REPARO DE HARDWARE DE 6 HORAS	O técnico deve chegar ao local em até 4 horas após o despacho e, em muitos casos, repara o hardware em até 6 horas após o despacho.	 Disponível 24 (vinte e quatro) horas por dia, 7 (sete) dias por semana, inclusive feriados. Para incidentes de Severidade 1, a Dell Technologies fará esforços comercialmente razoáveis para devolver o hardware à condição operacional em até 6 horas após o despacho Resposta em 4 horas e reparo em 6 horas após o despacho. Aplica-se somente a falhas ou reparos do Produto compatível coberto. O Suporte de software não está no escopo. Disponível para clientes a uma distância de até 80 quilômetros (50 milhas) do HUB de suporte designado da Dell Technologies. O Cliente deve ter uma versão compatível ativada e mantida do software de gateway de conexão segura.



RECURSO DO SERVIÇO	DESCRIÇÃO	PROSUPPORT PLUS — DETALHES DA COBERTURA
PLATAFORMAS AIOPS DO PROSUPPORT	AlOps é a inteligência artificial (IA) para as operações de TI. Ela se refere ao uso estratégico de tecnologias de IA, aprendizado de máquina (ML) e raciocínio de máquina (MR) que simplificam e facilitam processos e otimizam o uso dos recursos de TI do cliente.	Incluso. As plataformas AIOps do ProSupport incluem o APEX AIOps Infrastructure Availability, o TechDirect e o MyService 360, que são todos ativados pelo gateway de conexão segura da Dell Technologies, e oferecem benefícios não limitados a: • Detecção proativa de problemas e criação de casos • Detecção preditiva de falhas de hardware • Criação de casos de autoatendimento • Despacho de peças de autoatendimento • Avisos de segurança da Dell • Avaliação de segurança cibernética do APEX AIOps
		Infrastructure Availability O APEX AlOps Infrastructure Availability é um aplicativo de AlOps baseado em nuvem que fornece monitoramento e solução de problemas simples e proativos de sua infraestrutura de TI da Dell. Ele utiliza o aprendizado de máquina para monitorar e medir proativamente a integridade geral de servidores, armazenamento e dispositivos convergentes, hiperconvergentes, de proteção de dados e de rede por meio de análise preditiva, inteligente e abrangente. A análise preditiva da capacidade e do desempenho dos componentes, como unidades de estado sólido e memória, é habilitada por meio do uso do APEX AlOps Infrastructure Availability. O APEX AlOps Infrastructure Availability está disponível sem cobrança adicional para produtos com um contrato ProSupport (ou superior) válido. O APEX AlOps Infrastructure Availability é hospedado na nuvem privada da Dell Technologies, proporcionando a cada cliente um portal seguro e independente e garante que os clientes só poderão ver seu próprio ambiente.
		 Exibição proativa de incidentes e riscos críticos em tempo real Recomendações claras e prescritivas para simplificar a ação e o planejamento Tendências históricas e lógica analítica de dados Gerenciamento da integridade dos serviços para proteção de dados, armazenamento e sistemas convergentes
		O TechDirect permite o despacho automático de peças pelo cliente. O Gateway de conexão segura (SCG) é uma tecnologia de monitoramento empresarial fornecida como um equipamento e um aplicativo independente. Ele monitora seus dispositivos e detecta proativamente problemas de hardware que podem ocorrer. Dependendo do contrato de serviço, ele automatiza também a criação de solicitações de suporte para problemas detectados nos dispositivos monitorados.

Entre em contato com seu representante da Dell Technologies para saber mais sobre os produtos cobertos



PEÇAS DE REPOSIÇÃO	A Dell Technologies fornecerá peças de substituição quando considerar necessário.	Incluso. O objetivo da entrega de peças de reposição é uma resposta de serviço em quatro horas (ou seja, no mesmo dia útil) para peças essenciais depois que a Dell Technologies considerar necessária uma entrega de peça de reposição.
		As localizações das peças entregues em quatro (4) horas mantêm em estoque componentes do sistema, conforme determinado pela Dell Technologies. As peças não essenciais podem ser enviadas por meio de entrega noturna. As peças consideradas não essenciais incluem, entre outras: bordas, chassi mecânico, placas de proteção de disco rígido, kits de trilhos e acessórios de gerenciamento de cabos. As peças que podem ser consideradas críticas são: placas-mãe, CPUs, alguns módulos de memória e unidade de disco rígido que podem afetar o array e o ambiente de produção.
		Os horários limite do país local podem afetar a entrega no mesmo/ próximo dia útil local das peças de substituição não essenciais.
		Atualmente, a Dell tem peças em estoque em vários locais do mundo. As peças selecionadas podem não estar estocadas na localização mais próxima do cliente. Se uma peça necessária para reparar Produtos com suporte não estiver disponível em uma instalação Dell próxima ao local do Cliente e precise ser transferida de outra instalação, o envio ocorrerá durante a noite ou dentro do prazo mais breve que for prático e comercialmente razoável. As localizações das peças entregues em quatro (4) horas mantêm em estoque componentes do sistema, conforme determinado pela Dell. Para receber peças em quatro (4) horas, o Cliente deve estar localizado na área de cobertura determinada pela Dell. Para obter mais detalhes sobre limitações geográficas, consulte Termos e Condições Suplementares abaixo. Os tempos de resposta de entrega de peças mencionados acima são aplicáveis apenas para estocar peças de acordo com as configurações padrão. O estoque de peças de serviço para novos produtos ou configurações fora do padrão pode exigir um prazo mais longo, que levará trinta (30) dias.
		A instalação de todas as peças de substituição é feita pela Dell Technologies como parte da resposta no local. No entanto, o Cliente tem a opção de fazer a instalação de unidades substituíveis pelo cliente (CRUs). Consulte a <u>Tabela de garantia e manutenção de</u> <u>produtos da Dell Technologies</u> para obter uma lista das peças designadas como CRUs para Equipamento específico ou entre em contato com a Dell Technologies para obter mais detalhes.
		Se a Dell Technologies instalar a peça de substituição, ela providenciará a devolução da peça substituída a uma instalação da Dell Technologies. Se o Cliente instalar a CRU, ele será o responsável pela devolução da CRU substituída à instalação designada pela Dell Technologies. Se um cliente precisar de assistência para substituir uma CRU, a Dell Technologies poderá ajudar remotamente e/ou enviar um técnico no local para ajudar na substituição.
		Caso o técnico da Dell conclua, durante o diagnóstico, que o reparo pode ser realizado com uma peça designada como CRU ou se o Cliente optar por despachar automaticamente uma peça designada como CRU, a Dell enviará essa peça diretamente ao Cliente.
		Se o técnico da Dell determinar que o Produto com Suporte deve ser substituído por inteiro, a Dell Technologies reserva-se o direito de enviar ao Cliente uma unidade de substituição inteira. As substituições de unidades inteiras podem não estar disponíveis em estoque para tempos de resposta no mesmo dia, e pode haver prazos estendidos



		para a chegada de uma substituição de unidade inteira em sua localização, dependendo de onde você está localizado e do tipo
		de Produto que está sendo substituído.
SUBSTITUIÇÃO	Se, antes de atingir seu Nível de	Incluída para Produtos de armazenamento e infraestrutura
PROATIVA DE	Resistência, uma unidade de estado	hiperconvergente/convergente.
UNIDADE SSD	Resistência (conforme determinado	O objetivo de resposta baseja-se nos recursos de servicos aplicáveis
	pela Dell) ou superior, o Cliente será	de Entrega de peças de substituição e Resposta no local, detalhados
	elegível para receber uma unidade	acima. O cliente deve ativar e manter as versões atualmente
	de estado sólido de substituição.	compatíveis do software de monitoramento e suporte remoto de
	"Nivel de resistencia" significa a vida	II (Implementado como um gateway de conexao segura) durante
	"Limite de nível de resistência"	conectividade, conforme aplicavel, A alivação do software de
	é o ponto na vida útil da SSD	esses outros recursos de serviço de renovação.
	(conforme determinado pela Dell)	
	no qual a unidade se torna elegível	Unidades pré-criadas não são elegíveis para Substituição proativa
	para substituição, por exemplo, ao atingir 95% do pível de resistência	da unidade SSD pela Dell Technologies.
	Os limites de nível de resistência	
	variarão.	
DIREITOS	A Dell Technologies oferece os	Incluso.
A NOVAS	direitos as novas versoes de	
VERSÕES	sonware comorne as disponibiliza.	
DE SOFTWARE		
	A Dell Technologies realiza	Software do ambiente operacional (OF) do Equipamento
	a instalação remota das	
NOVAS VERSUES	novas Versões de software.	Incluído apenas no equipamento de armazenamento quando o
DE SUFTWARE		software associado do ambiente operacional incorporado está coberto
		por uma garantia da Dell ou por um contrato de manutenção em vigor
		definido como programação e/ou microcódigo do software de interface
		do usuário necessário para permitir a administração, o controle e a
		execução das funções básicas do equipamento e sem os quais o
		equipamento não pode operar.
		O cliente tem direito à instalação remota das atualizações de
		software do OE com uma versão compatível ativada e mantida do
		software do gateway de conexão segura.
		Consulte a Tabela de garantia e manutenção de produtos para
		produtos elegíveis.
		Quitros softwares (não OE)
		O Cliente realiza a instalação de novas Versões de Software, a menos
		que seja considerada necessária pela Dell Technologies.
MONITORAMENTO	Certos Produtos acionam um contato	Incluso para Produtos que têm SCG ou outras ferramentas de
REMOTO	automático e independente para	remoto disponível na Dell Technologies
E REPARO 24X7	fornecer dados que auxiliem	remoto disponíventa Deli reonnologica.
	a Dell Technologies a determinar	Veja os detalhes sobre a ferramenta SCG acima na seção Plataforma
	o problema.	de AIOPs.
	A Dell Technologies acessa os	Quando a Dell Technologies for notificada sobre um problema, os
	Produtos remotamente, caso	mesmos objetivos de resposta do Suporte técnico global e da
	necessario, para obter diagnosticos adicionais e dar suporte remoto	Resposta no local serão aplicáveis, conforme descrito anteriormente.
	actionale e dai capone remoto.	

Dell Technologies ProSupport Plus for Infrastructure | v11 | May 2024



SERVICE ACCOUNT MANAGER ("SAM") O SAM atribuído do ProSupport Plus for Infrastructure é um recurso remoto que oferece uma ampla variedade de recursos e funcionalidades de gerenciamento de sistema, ambiente e conta, criados para reduzir o tempo de inatividade e melhorar a experiência geral de suporte da Dell Technologies.

Estão incluídos no Serviço: Assistência de integração:

- Verificando a exatidão das informações relevantes de suporte ao cliente, como nome da conta, endereço etc.
- Fornecendo transferências de conhecimento, tais como entrar em contato com a Dell Technologies para abrir chamados e o uso de ferramentas e tecnologias de suporte da Dell Technologies
- Designando o cronograma para entregas do SAM, como relatórios e análises de serviço

<u>Relatórios mensais</u>: Geração de relatórios e recomendações sobre sistemas qualificados, incluindo:

- Resumo das chamados abertos e fechados por mês;
- Verificação de versões de software do sistema atualmente instaladas com base em recomendações de código de destino; e
- Status do contrato, incluindo datas de início/fim e outros detalhes básicos do contrato.

Para permitir totalmente a geração de relatórios mensais, as tecnologias de conectividade da Dell Technologies, como o gateway de conexão segura, devem ser instaladas com as opções apropriadas de coleta de log ativadas.

Análise de serviço: O SAM fornece uma análise de serviço dos detalhes no relatório de serviço. Prazo, cronograma e outros temas a serem analisados serão determinados entre o SAM e o Cliente durante a integração.

Manutenção do sistema: Para ativos qualificados, o SAM auxiliará o Cliente na coordenação da entrega de eventos de manutenção do sistema dentro da janela de manutenção do Cliente. Veja abaixo informações adicionais.

<u>Suporte de escalonamento da</u> <u>Dell Technologies</u>: atuar como elo de Serviços para coordenar os recursos necessários e resolver questões individuais de Severidade 1 ou problemas mais sistêmicos. Incluído nos Produtos cobertos pelo serviço ProSupport Plus for Infrastructure ou no contrato de manutenção em vigor durante o horário comercial local normal da Dell Technologies, que pode variar de acordo com a região e o país, exceto em feriados locais e da Dell Technologies. Consulte os detalhes adicionais da cobertura abaixo.

A Dell Technologies é responsável por realizar apenas as atividades e tarefas do SAM expressamente especificadas neste documento. Todas as outras tarefas, atividades e serviços estão fora do escopo.



DEFINIÇÕES DE NÍVEL DE SEVERIDADE

SEVERIDADE 1 Crítica — incapaz de executar funções essenciais aos negócios e requer resposta imediata.

SEVERIDADE 2: Alta — capaz de executar funções de negócios, mas o desempenho/capacidades estão degradados ou gravemente limitados.

SEVERIDADE 3 Média/baixa - mínimo ou nenhum impacto nos negócios.

Informações adicionais sobre o Suporte essencial

A Dell se reserva o direito de recusar o serviço se a Dell Technologies acreditar que o Cliente está usando incorretamente ou em demasia o recurso On-Demand Onsite Diagnosis para problemas críticos (por exemplo, quando o cliente tiver uma equipe disponível para solucionar o problema no local ou as solicitações do Cliente para as visitas de diagnóstico no local excederem significativamente as taxas de falha padrão, em decorrência dos exemplos a seguir, dos componentes e sistemas envolvidos). Se a Dell Technologies determinar (a critério exclusivo da Dell Technologies) que um Cliente está utilizando abusivamente o Serviço, a Dell Technologies se reservará o direito de negar o serviço.

Exclusões

As seguintes atividades não estão incluídas no escopo desta Descrição de serviço:

- Desinstalação, reinstalação ou configuração de produtos, software ou aplicativos
- Remoção de Produto desinstalado do local do Cliente
- O software de servidor/armazenamento/rede não é classificado como equipamento
- Solução de problemas de software do ambiente operacional que não reabilitarão o produto a um estado funcional (por exemplo, a consultoria, o ajuste de desempenho, a configuração, a criação de scripts ou a análise comparativa não está incluído)
- Serviços necessários devido à falta de manutenção do software e dos Produtos compatíveis com o suporte em qualquer nível mínimo de versão especificado, conforme estabelecido no <u>Documento de código de referência</u>.
- Serviços necessários devido à não execução de qualquer correção, reparo, patch ou modificação do sistema fornecidos pela Dell Technologies ou à falha do Cliente em tomar medidas de contenção previamente recomendadas pela Dell Technologies, como avisos de segurança ou atualizações de correção essenciais comunicados e que o cliente não implementou
- Serviços que, segundo a Dell Technologies, são necessários devido ao tratamento ou uso inadequado dos produtos ou
 equipamentos
- Personalização do servidor ou do dispositivo de armazenamento do Cliente, exceto quando expressamente declarado nesta Descrição de serviço
- Qualquer recuperação ou transferência de dados ou aplicativos
- Serviço de garantia ou suporte para sistemas, software ou componentes adicionais que não sejam da Dell Technologies
- Serviços que, segundo a Dell Technologies, são necessários devido a tentativas não autorizadas por pessoal terceirizado de instalar, reparar, manter ou modificar hardware, firmware ou software
- Instalação da impressora de rede ou mapeamento do compartilhamento de arquivos de rede
- Configuração de qualquer tipo para servidor, armazenamento, rede ou roteador
- Serviços de rede, incluindo conexão de um sistema a uma rede (diferente de uma LAN Ethernet)
- Qualquer atividade não estabelecida especificamente nesta Descrição de serviço.

Esta Descrição de serviço não concede ao Cliente nenhuma garantia adicional às garantias fornecidas segundo os termos do contrato principal de serviços ou do Contrato, conforme aplicável.



Responsabilidades do Cliente pelo recurso do serviço do SAM

O fornecimento por parte da Dell Technologies do recurso de serviço SAM detalhado acima dependerá do cumprimento do Cliente das seguintes responsabilidades:

- Disponibilizar as janelas de manutenção do sistema apropriadas para o SAM conforme considerado necessário pela Dell Technologies.
- Garantia de que todos os requisitos ambientais, técnicos e operacionais sejam atendidos.
- Fornecer ao SAM acesso oportuno a (a) pelo menos um contato técnico com responsabilidades de administração do sistema e privilégios apropriados de acesso ao sistema/informações, e (b) especialistas no assunto, sistemas e redes aplicáveis (incluindo, entre outros, acesso a sistemas remotos/rede) conforme considerado necessário pela Dell Technologies.
- Assumir toda a responsabilidade pela conectividade de rede, pelo desempenho e pelos problemas de configuração.
- Verificar se as localizações dos Equipamentos estão preparadas antes do início do ProSupport Plus for Infrastructure.
- Garantir que o produto esteja conectado ao gateway de conexão segura com as opções apropriadas de coleta de log ativadas.

Outras informações importantes sobre o recurso de serviço do SAM

- O serviço do SAM está disponível durante o horário comercial normal. O horário comercial é definido pelo local onde o SAM reside e pode variar de acordo com a região e o país. A critério da Dell Technologies e quando ela considerar necessário, os serviços do SAM podem ser realizados no local.
- O suporte depois do expediente pode ser fornecido por outros recursos da Dell Technologies a critério da Dell Technologies.
- O local do SAM será atribuído por acordo mútuo durante a integração, com base na área de serviço preferencial do cliente e na disponibilidade da equipe da Dell.

MANUTENÇÃO PROATIVA DO SISTEMA PROSUPPORT PLUS FOR INFRASTRUCTURE

A Manutenção do sistema do ProSupport Plus fornece aos clientes da Dell Technologies eventos de manutenção remota proativa e planejada que ocorrem durante a vigência do contrato de serviço em dispositivos cobertos pelo ProSupport Plus for Infrastructure e monitorados com a tecnologia de conectividade aplicável da Dell Technologies, como o gateway de conexão segura, conforme aplicável. Esses eventos de manutenção do sistema ajudam a manter o desempenho e podem reduzir a probabilidade de incidentes futuros devidos a versões incompatíveis de hardware, software, BIOS e firmware. Os eventos de manutenção de sistema proativa e planejada são coordenados entre os clientes, o SAM e o pessoal de suporte da Dell Technologies. A entrega desses eventos geralmente está disponível 24x7x365, mas pode estar sujeita à mútua disponibilidade de recursos do cliente e da Dell Technologies. A Dell Technologies recomenda que a manutenção de sistema proativa e planejada ocorra duas vezes por ano. Certos produtos podem ter limitações no número de vezes que a manutenção de sistema proativa e planejada pode ser realizada por ano. Consulte o representante de vendas ou o SAM designado para obter uma lista de Produtos com Suporte e saber guais são as limitações aplicáveis.

Não incluídos na manutenção do sistema do ProSupport Plus for Infrastructure

- Atualizações em dispositivos interconectados não cobertos por um contrato de suporte atual do ProSupport Plus for Infrastructure.
- Atualizações em qualquer software sem o direito correspondente a tais atualizações de acordo com um contrato de suporte de software apropriado por parte da Dell Technologies ou de terceiros para alguns Produtos de Terceiros.
- Upgrades de sistema operacional e criação de patches de hypervisor ou outro suporte de engenharia ou de desenvolvimento de software relacionado.
- Criação de patches de aplicativo.
- Prestação de serviços de manutenção no local.
- Desinstalação ou instalação de hardware adicional, ou tarefas de configuração.
- Instalação ou configuração de software não especificamente listado nesta descrição de serviço.
- Ajuste de desempenho do aplicativo.
- Remoção ou identificação de vírus, spyware, ou malware.
- Todas as outras atualizações ou outras atividades não especificamente documentadas dentro desta Descrição do serviço.



Informações adicionais importantes sobre a manutenção do sistema do ProSupport Plus for Infrastructure

- Durante o evento de manutenção, upgrades podem causar uma perda temporária de conectividade com outros dispositivos conectados.
- Após a conclusão da atualização, os dispositivos conectados talvez precisem ser reiniciados e a conectividade verificada.
- Os sistemas a serem submetidos a upgrade devem ser disponibilizados para a Dell Technologies ou para os agentes autorizados pela Dell Technologies durante a janela de manutenção acordada.
- Dependendo dos sistemas a serem atualizados, pode ser necessário disponibilizar um sistema ou recurso de gerenciamento de sistema adicional.
- Dependendo dos sistemas a serem submetidos a upgrade, pode ser necessário fornecer direitos de administração apropriados para o dispositivo à Dell Technologies ou aos agentes autorizados da Dell Technologies.
- O cliente é responsável por ter e manter todos os requisitos de licença relativos às atualizações de hardware e software.
- Caso essa atualização do Software em um Produto qualificado possa reduzir ou afetar o desempenho de outro Produto não qualificado, a Dell Technologies, em consulta ao cliente, pode optar por não dar andamento à atividade de manutenção do sistema até que a situação seja resolvida.
- Para permitir totalmente a Manutenção de Sistema do ProSupport Plus for Infrastructure, a tecnologia de conectividade aplicável, como o gateway de conexão segura, deve ser instalada com as opções de coleta de log ativadas.
- A instalação de novas Versões de software para os sistemas de armazenamento de avançados e os sistemas de infraestrutura convergente e hiperconvergente da Dell Technologies determinadas durante a manutenção do sistema, inclusive, entre outras, versões de software publicadas nas matrizes de configuração de interoperabilidade aplicáveis (a Matriz de suporte simples ou a Matriz de certificação de versão da Dell Technologies), pode exigir a compra de um projeto de serviços separado da Dell Technologies para obter mais informações.

Entrega de relatórios a revendedores autorizados da Dell Technologies para Clientes usuários finais que compram de revendedores autorizados da Dell Technologies

O relatório do SAM do ProSupport Plus será fornecido pela Dell Technologies apenas ao **revendedor ou distribuidor autorizado da Dell Technologies (conforme aplicável) identificado na cotação da Dell Technologies (coletivamente, o "Revendedor"). O Revendedor reconhece e concorda que obteve consentimento apropriado dos clientes usuários finais do Revendedor ("Usuários Finais") para receber os Relatórios do SAM do cliente Usuário Final**. A não ser que tenha sido notificado em contrário pelo Revendedor, a Dell Technologies distribuirá os relatórios do SAM do ProSupport Plus para o endereço e as informações de contato do Revendedor fornecidos à Dell Technologies quando o Serviço foi adquirido ou para as informações de contato contidas nos registros atuais de vendas e serviços da Dell Technologies como determinado por ela. Os relatórios do SAM fornecidos ao Revendedor não serão categorizados por/para Usuários Finais específicos. Opções de relatórios personalizados podem estar disponíveis para compra separada por um custo adicional.

ASSISTÊNCIA COLABORATIVA

Se o Cliente abrir um chamado e a Dell Technologies determinar que o problema surge com um produto de fornecedor terceirizado elegível comumente utilizado em conjunto com os Produtos cobertos por uma garantia ou por um contrato de manutenção atual da Dell Technologies, a Dell Technologies fará esforços para fornecer Assistência colaborativa, em que a Dell Technologies: (i) serve como ponto único de contato até que os problemas sejam isolados; (ii) entra em contato com o fornecedor terceirizado; (iii) fornece documentação do problema; e (iv) continua a monitorar o problema e obtém o status e planos de resolução do fornecedor (quando razoavelmente possível).

Para ter direito à Assistência colaborativa, o Cliente deve ter os contratos de suporte ativos e os direitos apropriados diretamente com o respectivo fornecedor terceirizado e a Dell Technologies ou com um revendedor autorizado Dell Technologies. Depois que o problema é isolado e informado, o fornecedor terceirizado torna-se o único responsável por oferecer todo o suporte, seja ele técnico ou não, em conexão com a resolução do problema do Cliente. A Dell Technologies NÃO SE RESPONSABILIZA PELO DESEMPENHO DE PRODUTOS OU SERVIÇOS DE OUTROS FORNECEDORES. Uma lista de parceiros de Assistência colaborativa pode ser encontrada na Lista de assistência colaborativa. Observe que os produtos cobertos pelo suporte de terceiros podem ser alterados a qualquer momento, sem qualquer aviso prévio aos Clientes.



SUPORTE DE SOFTWARE DO SISTEMA DA DELL TECHNOLOGIES

O suporte de software da Dell Technologies incluído no ProSupport Plus for Infrastructure destina-se a alguns Produtos de terceiros, incluindo determinados aplicativos de usuário final, sistemas operacionais, hypervisores e firmware, quando esses Produtos de terceiros 1) estão instalados e funcionando e são usados com os Produtos no momento em que o suporte é solicitado, 2) estão cobertos por um termo de servico de suporte e manutenção do ProSupport Plus for Infrastructure e 3) têm os direitos e contratos de suporte ativos apropriados diretamente com o respectivo editor de Produtos de terceiros. Esse nível de suporte é fornecido no Equipamento qualificado do ProSupport Plus for Infrastructure, independentemente de como o software qualificado foi adquirido e licenciado, mas o Cliente é responsável por garantir que esse software qualificado tenha sido adquirido e licenciado corretamente de acordo com o editor. O Cliente é o único responsável por corrigir quaisquer problemas com licenças e compras de software qualificado para poder receber esses Serviços a qualquer momento durante o período de cobertura. Os aplicativos de software elegíveis podem ser encontrados na Lista abrangente de suporte de software. Os produtos suportados de terceiros podem ser alterados a qualquer momento, sem aviso aos Clientes. As situações que dão origem às perguntas do cliente devem ser reprodutíveis em um sistema único, que pode ser físico ou virtual. O cliente compreende e aceita que as soluções de alguns problemas que dão origem ao chamado do Cliente podem não ser fornecidas pelo editor do software em questão (incluindo, entre outras, situações em que o editor não está mais fornecendo suporte ou manutenção ao software em questão por qualquer motivo) ou podem exigir o suporte adicional do editor, inclusive a instalação de software adicional ou outras alterações nos produtos. O cliente aceita que, nessas situações nas quais o editor do software em questão não fornece uma resolução, a obrigação da Dell Technologies de fornecer suporte ao cliente será plenamente cumprida.

Termos e condições adicionais aplicáveis aos usuários finais que compram Produtos de um OEM

O "OEM" é um revendedor que vende os Produtos compatíveis como um fabricante de equipamento original que está adquirindo os Produtos e Serviços da Dell Technologies do grupo de negócios OEM Solutions (ou seu sucessor) para um projeto de OEM. O OEM normalmente incorpora ou agrupa esses Produtos da Dell Technologies com hardware, software ou outra propriedade intelectual pertencente ao Cliente OEM, resultando em um sistema ou solução especializado com a funcionalidade específica da tarefa ou do setor (esse sistema ou solução sendo uma "Solução de OEM") e revende essa Solução de OEM com sua própria marca. Com relação a OEMs, o termo "Produtos com Suporte" inclui os Produtos com Suporte da Dell Technologies que são fornecidos sem a marca Dell Technologies (ou seja, sistema pronto de OEM sem marca), e "Usuário final" significa você ou qualquer entidade que adquire uma Solução de OEM para sua própria utilização final e não para revenda, distribuição ou sublicenciamento a terceiros. É responsabilidade do OEM oferecer ao Usuário Final solução de problemas de primeiro nível. Um diagnóstico inicial apropriado de melhor esforço deverá ser realizado pelo OEM antes de a chamada prosseguir para a Dell Technologies. Esse OEM permanece com a responsabilidade de fornecer a solução de problemas inicial, mesmo quando o Usuário final contata a Dell Technologies para solicitar o serviço. Além disso, se um Usuário final contatar a Dell Technologies para solicitar o serviço. Além disso, se um Usuário final contatar seu OEM, a Dell Technologies pedirá que o Usuário final contate o OEM para receber uma solução de problemas de primeiro nível antes de entrar em contato com a Dell Technologies.

Dell Technologies ProSupport Plus for Infrastructure sobre peças não padrão em produtos personalizados de servidor

Os reparos e substituições de peças não padrão ou únicas ("Serviços de Suporte de Componente Não Padrão") são um serviço de substituição de valor agregado que complementa a garantia do Produto PowerEdge do Cliente em componentes padrão da Dell Technologies em uma configuração padrão e que exigem substituição devido a defeitos de fabricação ou de material ("Reparos na Garantia"). Firmware/software com a marca Dell Technologies NÃO estão disponíveis para "Componentes Não Padrão", e o Cliente deve usar utilitários fornecidos pelo fabricante para monitorar e/ou atualizar o componente. O Cliente trabalhará diretamente com o fabricante para resolver qualquer problema de qualidade relativo a software/firmware, utilitários e hardware. A Dell Technologies prestará Serviços de Suporte Não Padrão para substituir peças não padrão ou únicas que o Cliente preveja e garanta que estarão disponíveis como definido acima e desde que o Cliente tenha feito os arranjos correspondentes para ajudar a Dell Technologies no processo de pedidos para estoque de serviço a fim de facilitar a atividade de reparo. Desde que o Cliente tenha uma previsão exata das necessidades de estoque, a Dell Technologies trocará a peca que apresente defeito de acordo com o tempo de resposta aplicável do Cliente para Reparos na Garantia e instalará a peça de substituição no Produto do Cliente, mas o Cliente reconhece e concorda que a Dell Technologies não é responsável perante ele por garantir a disponibilidade das peças. As peças e a resposta em campo no mesmo dia (por exemplo, 4 horas) podem não estar disponíveis para substituição de componentes "não padrão". Nesses casos, a Dell Technologies definirá como padrão o serviço no próximo dia útil (ou seja, após o diagnóstico e a solução do problema por telefone, uma peça ou um técnico geralmente será despachado no próximo dia útil). As peças de substituição podem ser novas ou reformuladas conforme permitido pelas leis locais, e a realização de reparos e substituições de Serviços de Suporte de Componente Não Padrão pode exigir que a Dell Technologies utilize garantia e/ou serviços de manutenção de um fabricante/editor terceirizado, e o Cliente concorda em prestar assistência à Dell Technologies e em fornecer todo o material solicitado por qualquer fabricante ou editor terceirizado para facilitar a utilização da garantia e/ou dos serviços de manutenção do terceiro correspondente.



Os testes de engenharia da configuração resultante realizados pela Dell Technologies e baseados em uma declaração de trabalho SOW separada, por exemplo, testes executados depois da instalação das peças não padrão ou exclusivas para uma configuração que utiliza o software solicitado pelo cliente, são uma atividade pontual a ser realizada uma vez, e não continuamente, e os serviços de suporte de componente não padrão estão disponíveis apenas para a configuração específica conforme definida pelo cliente e testada pela Dell Technologies. A Dell Technologies comunicará a configuração de hardware exata testada, inclusive os níveis de firmware. Depois que os testes de engenharia forem concluídos, a Dell Technologies fornecerá os resultados por meio de relatórios com a indicação de Aprovado/Reprovado. A Dell Technologies envidará esforços comercialmente razoáveis para dar suporte ao reconhecimento e à operação do componente não padrão no Produto da Dell Technologies, no entanto, a modificação de utilitários padrão da Dell Technologies (inclusive BIOS, IDRAC e software de conectividade) não será aceita. O Cliente será responsável por trabalhar diretamente com o fabricante para resolver qualquer problema do componente não padrão que surja durante o teste de engenharia (inclusive problemas de qualidade, software, firmware ou especificações/limitações de hardware). Os testes de engenharia adicionais da Dell Technologies depois que o Cliente tiver recebido um relatório com a indicação de APROVADO exigirão uma nova SOW e tarifas de engenharia não recorrentes associadas, inclusive em qualquer teste de engenharia solicitado em conexão com

Responsabilidades Gerais do Cliente

Autoridade para Conceder Acesso. O Cliente declara e garante que obteve permissão tanto para o Cliente quanto para a Dell Technologies de acessar e usar, seja remota ou pessoalmente, o software, o hardware, os sistemas licenciados ou pertencentes ao Cliente e os dados neles contidos, além de todos os componentes de hardware e software neles incluídos, para a finalidade de prestar esses Serviços. Se o Cliente ainda não tiver essa permissão, caberá a ele obtê-la, às suas custas, antes de solicitar que a Dell Technologies preste esses serviços.

Não solicitação. Quando permitido pela lei, o cliente não solicitará – sem a aprovação prévia por escrito da Dell Technologies, por um período de dois anos a partir da data indicada no seu Formulário de pedido – direta ou indiretamente a contratação de qualquer funcionário da Dell Technologies com quem ele tenha entrado em contato devido à prestação do Serviço pela Dell Technologies, desde que, no entanto, anúncios gerais e outras formas amplas de contratação não constituam solicitação direta ou indireta nos termos deste instrumento e você tenha permissão para propor a contratação de qualquer funcionário que tenha sido demitido ou tenha pedido demissão do emprego na Dell Technologies antes do início das conversas sobre emprego com você.

Cooperação do Cliente. O Cliente compreende que, sem cooperação rápida e adequada, a Dell Technologies não conseguirá realizar o Serviço ou, se realizado, o Serviço poderá ser significativamente alterado ou atrasado. Sendo assim, de maneira imediata e aceitável, o Cliente cooperará integralmente com a Dell Technologies, conforme necessário, para que ela possa realizar o Serviço. Se o Cliente não cooperar de forma razoavelmente adequada, de acordo com o disposto acima, a Dell Technologies não será responsável por nenhuma falha na realização do Serviço, e o Cliente não terá direito a reembolso.

Obrigações no Local. Quando os Serviços exigirem trabalho no local, o Cliente fornecerá (sem custo para a Dell Technologies) acesso livre, seguro e suficiente às instalações e ao ambiente do Cliente, incluindo amplo espaço de trabalho, eletricidade, equipamento de segurança (se aplicável) e linha telefônica local. O Cliente deverá fornecer também (sem custo para a Dell Technologies) um monitor ou uma tela, um mouse (ou dispositivo de apontamento eletrônico) e um teclado, caso o sistema já não inclua esses itens.

Backup dos Dados. O Cliente fará um backup de todos os dados, software e programas existentes em todos os sistemas afetados antes e durante a prestação deste Serviço. O Cliente deve fazer cópias de backup regulares dos dados armazenados em todos os sistemas afetados, como precaução contra possíveis falhas, alterações ou perdas de dados. A Dell Technologies não será responsável pela restauração ou reinstalação de nenhum programa nem de dados. Exceto quando proibido pelas leis locais aplicáveis, a Dell Technologies não terá responsabilidade por perda de dados em relação a:

- 1. Quaisquer informações confidenciais, de propriedade exclusiva ou pessoais;
- 2. Dados, programas ou software perdidos ou corrompidos;
- 3. Mídia removível perdida ou corrompida;
- 4. Perda do uso de um sistema ou rede; e/ou
- 5. Qualquer ato ou omissão, inclusive negligência, da Dell Technologies ou de um prestador de serviços terceirizado.

Garantias de Terceiros. Estes serviços podem exigir que a Dell Technologies tenha acesso a hardware ou software que não tenha sido produzido ou vendido por ela. As garantias de alguns fabricantes podem ser anuladas se a Dell Technologies ou outra pessoa que não o fabricante trabalhar no hardware ou no software. O Cliente garantirá que a execução dos serviços pela Dell Technologies não afetará tais garantias ou, se isso ocorrer, que o resultado será aceitável para o Cliente. A Dell Technologies não se responsabiliza por garantias de terceiros nem por efeitos que os serviços possam ter nessas garantias.



Manter as versões do Software e Serviço. O Cliente deverá manter o software e os Produtos compatíveis com o suporte, de acordo com as os níveis de versão mínimos especificados pela Dell Technologies no Documento de código de referência. O Cliente deverá também garantir a instalação dos níveis de versão mínimos de software e firmware nas peças de reposição, patches, atualizações de software ou versões subsequentes, conforme indicado pela Dell, para manter os direitos dos Produtos compatíveis com o suporte à realização do Serviço. A Dell Technologies reserva-se o direito, a seu critério exclusivo, de negar suporte a quaisquer software e Produtos compatíveis com o suporte que não atendam aos níveis mínimos de versão especificados pela Dell Technologies, conforme especificado no Documento de código de referência.

Termos e Condições dos Serviços

Esta Descrição de Serviço é acordada entre você, o Cliente ("você" ou "Cliente") e a Dell Technologies. Este Serviço está sujeito e é regido pelo Acordo do Cliente com a Dell Technologies.

Os produtos ou serviços obtidos de qualquer revendedor Dell Technologies serão regidos exclusivamente pelo contrato firmado entre o comprador e o revendedor. Esse contrato pode fornecer termos iguais aos deste documento ou dos termos on-line abaixo. O revendedor pode firmar acordos com a Dell Technologies para realizar serviços de garantia e/ou manutenção para o comprador em nome do revendedor. Os clientes e os revendedores que realizam serviços de garantia e/ou manutenção ou serviços profissionais devem ser adequadamente treinados e certificados. A realização de qualquer serviço por Clientes, revendedores ou terceiros não treinados/não certificados pode resultar em tarifas adicionais se o suporte da Dell Technologies for necessário em resposta à prestação de serviços por terceiros. Entre em contato com o revendedor ou o representante de vendas local da Dell Technologies para obter mais informações sobre a prestação de serviços de garantia e manutenção da Dell Technologies para Produtos adquiridos de um revendedor.

Na ausência de tal acordo autorizando explicitamente este Serviço e dependendo da localização do Cliente, este Serviço estará sujeito e será regido pelos Termos comerciais de venda da Dell ou pelo contrato de revenda ao qual a tabela a seguir fizer referência. Consulte a tabela abaixo, que mostra o URL aplicável à localização do Cliente onde seu contrato está disponível. As partes confirmam que leram e concordam com o cumprimento destes termos on-line.

Lecelizeeãe	- Termos e condições aplicáveis à compra dos Serviços		
- Localização do Cliente	 Clientes que compram os Serviços diretamente 	 Clientes que compram os Serviços por meio de um revendedor Authorized 	
- Estados Unidos	- <u>Dell.com/CTS</u>	- <u>Dell.com/CTS</u>	
- Canadá	- <u>Dell.ca/terms</u> (inglês) <u>Dell.ca/conditions</u> (francês do Canadá)	- <u>Dell.ca/terms</u> (inglês) <u>Dell.ca/conditions</u> (francês do Canadá)	
- Países da América Latina e do Caribe	Site <u>Dell.com</u> local específico do país ou <u>Dell.com/servicedescriptions</u> .*	As Descrições de serviço e outros documentos de serviço da Dell Technologies que você possa receber do vendedor não constituirão um contrato entre você e a Dell Technologies, servindo apenas para descrever o conteúdo do Serviço que você está adquirindo desse vendedor, suas obrigações como destinatário do Serviço e os escopos e limitações de tal Serviço. Como consequência e neste contexto, qualquer referência a "Cliente" nesta Descrição de Serviço e em outros documentos de serviço da Dell Technologies será entendida como uma referência a você, enquanto as referências à Dell Technologies devem ser entendidas como uma referência à Dell Technologies como um provedor de serviços que está fornecendo o Serviço em nome do seu vendedor. Você não terá um relacionamento contratual direto com a Dell Technologies em relação ao Serviço descrito neste documento. Para evitar dúvidas, qualquer condição de pagamento ou contratual que seja, por natureza, diretamente relevante apenas entre um comprador e um vendedor não será aplicável a seu caso e seguirá o acordado entre você e seu vendedor.	
- Ásia- Pacífico e Japão	Site <u>Dell.com</u> local específico do país ou <u>Dell.com/servicedescriptions</u> .*	As Descrições de Serviço e outros documentos de serviço da Dell Technologies que você possa receber do vendedor não constituem um contrato entre você e a Dell Technologies, servindo apenas para descrever o conteúdo do Serviço que você está adquirindo desse vendedor, suas obrigações como destinatário do Serviço e os escopos e limitações de tal Serviço. Como consequência e neste contexto, gualguer referência a "Cliente"	



		nesta Descrição de Serviço e em outros documentos de serviço da Dell Technologies será entendida como uma referência a você, enquanto as referências à Dell Technologies devem ser entendidas como uma referência à Dell Technologies como um provedor de serviços que está fornecendo o Serviço em nome do seu vendedor. Você não terá um relacionamento contratual direto com a Dell Technologies em relação ao Serviço descrito neste documento. Para evitar dúvidas, qualquer condição de pagamento ou contratual que seja, por natureza, diretamente relevante apenas entre um comprador e um vendedor não será aplicável a seu caso e seguirá o acordado entre você e seu vendedor.
- Europa, Oriente Médio e África	 Site <u>Dell.com</u> local específico do país ou <u>Dell.com/servicedescriptions</u>.* Além disso, os Clientes localizados na França, na Alemanha e no Reino Unido podem selecionar o URL aplicável abaixo: França: <u>Dell.fr/ConditionsGeneralesdeVente</u> Alemanha: <u>Dell.de/AGB</u> Reino Unido: <u>Dell.co.uk/terms</u> 	As Descrições de Serviço e outros documentos de serviço da Dell Technologies que você possa receber do vendedor não constituem um contrato entre você e a Dell Technologies, servindo apenas para descrever o conteúdo do Serviço que você está adquirindo desse vendedor, suas obrigações como destinatário do Serviço e os escopos e limitações de tal Serviço. Como consequência e neste contexto, qualquer referência a "Cliente" nesta Descrição de Serviço e em outros documentos de serviço da Dell Technologies será entendida como uma referência a você, enquanto as referências à Dell Technologies devem ser entendidas como uma referência à Dell Technologies como um provedor de serviços que está fornecendo o Serviço em nome do seu vendedor. Você não terá um relacionamento contratual direto com a Dell Technologies em relação ao Serviço descrito neste documento. Para evitar dúvidas, qualquer condição de pagamento ou contratual que seja por natureza diretamente relevante anenas
		entre um comprador e um vendedor não será aplicável a seu caso e seguirá o acordado entre você e seu vendedor.

* Os clientes podem acessar o site <u>Dell.com</u> local acessando <u>Dell.com</u> em um computador conectado à Internet em sua localidade ou escolhendo entre as opções do site "Choose a Region/Country" (Escolha uma Região/País) da Dell disponível em <u>Dell.com/content/public/choosecountry.aspx?c=us&l=en&s=gen</u>.

O cliente também concorda que, se renovar, modificar, estender ou continuar utilizando o Serviço além do período de vigência inicial, o Serviço estará sujeito à Descrição de Serviço em vigor na época e disponível para análise em <u>Dell.com/servicedescriptions</u>.

Ao fazer o pedido dos Serviços, receber a prestação dos Serviços, utilizar os Serviços ou o software associado ou clicar no botão/marcar a caixa "I Agree" (Eu concordo) ou similar em relação à sua compra no site Dell.com ou DellEMC.com, em uma interface da Internet ou um software da Dell Technologies, você concorda com o cumprimento desta Descrição de Serviço e com os acordos incorporados a ela por referência. Se você está firmando esta Descrição de Serviço em nome de uma empresa ou outra pessoa jurídica, você declara ter autoridade para vinculá-la a esta Descrição de Serviço e, nesse caso, "você" ou o "Cliente" deverá se subordinar à entidade em questão. Além de receberem esta Descrição de Serviço, os Clientes de determinados países também poderão ser solicitados a assinar um Formulário de pedido.



Termos e Condições Adicionais

Vigência do Serviço. Esta Descrição de Serviço se inicia na data registrada no Formulário de Pedido e continua até o final da vigência ("Vigência") indicada no Formulário de pedido. A quantidade de sistemas, licenças, instalações, implementações, endpoints ou usuários finais gerenciados para os quais o Cliente adquiriu um ou mais Serviços, o valor ou o preço e a Vigência aplicável a cada Serviço estão indicados no Formulário de Pedido do Cliente, conforme aplicável. Salvo se acordado por escrito de outra forma entre a Dell Technologies e o Cliente, as aquisições dos Serviços incluídos nesta Descrição de serviço serão exclusivamente para o uso interno do Cliente, e não para fins de revenda nem para agenciamento de serviços.

2. Informações adicionais importantes

- a) Reagendamento. Quaisquer alterações na programação após o agendamento deste Serviço devem ocorrer pelo menos 8 dias corridos antes da data agendada. Se o Cliente fizer um novo agendamento do serviço a sete dias ou menos da data agendada, será cobrada uma tarifa de novo agendamento, que não deverá exceder 25% do preço pago pelo Cliente para a prestação dos Serviços. Qualquer reagendamento do Serviço deverá ser confirmado pelo Cliente pelo menos 8 dias antes do início do Serviço.
- b) Pagamento do Hardware Adquirido com os Serviços. Salvo se acordado de outra forma por escrito, o pagamento do hardware, em nenhuma hipótese, dependerá da prestação ou do fornecimento dos Serviços comprados com ele.
- c) Limites do escopo dos Serviços. A Dell poderá se recusar a prestar o Serviço se, em sua opinião, a prestação do Serviço criar um risco não aceitável à Dell ou aos provedores de serviços da Dell ou se algum serviço solicitado estiver fora do escopo do Serviço. A Dell não se responsabiliza por nenhuma falha ou atraso na execução devido a causas fora de seu controle, inclusive a incapacidade de o Cliente cumprir com suas obrigações nos termos desta Descrição de serviço.
- d) Alterações no Escopo do Serviço. Salvo se acordado de outra forma por escrito com o Cliente, a Dell Technologies se reserva o direito de alterar o escopo dos Serviços em sessenta (60) dias antes do prévio aviso por escrito ao Cliente. Além das alterações causadas por editores e fabricantes de Produtos de Terceiros, o Cliente será notificado sobre qualquer alteração no escopo dos Serviços de Suporte da maneira declarada no Contrato atual entre a Dell Technologies e o Cliente.
- e) Privacidade. A Dell Technologies tratará todas as informações pessoais coletadas sob esta Descrição de serviço de acordo com a Declaração de Privacidade da Dell Technologies da jurisdição aplicável, todas as quais estão disponíveis em http://www.dell.com/localprivacy e cada uma delas é incorporada por referência.
- f) Serviços Opcionais. Serviços opcionais (inclusive suporte em pontos de necessidade, instalação, consultoria, serviços gerenciados, de suporte, profissionais ou de treinamento) podem estar disponíveis para compra com a Dell Technologies e variarão de acordo com a localização do Cliente. Os serviços opcionais podem exigir um contrato separado com a Dell Technologies. Na ausência de tal contrato, os serviços opcionais são fornecidos de acordo com esta Descrição de Serviço.
- g) Atribuição e Terceirização. A Dell Technologies pode terceirizar este Serviço e/ou atribuir esta Descrição de serviço a prestadores de serviço terceirizados qualificados, que prestarão o Serviço em nome da Dell Technologies.
- h) Cancelamento. A Dell Technologies pode cancelar este Serviço a qualquer momento durante a Vigência por qualquer um dos seguintes motivos:
 - a. O Cliente deixar de pagar o preço total ou parcial deste Serviço de acordo com os termos da fatura;
 - b. O Cliente for abusivo, ameaçador ou se recusar a cooperar com o analista assistente ou o técnico no local;
 - c. O Cliente deixar de respeitar os termos e condições estabelecidos nesta Descrição de serviço;
 - O Cliente solicita a substituição de componentes que excedem materialmente as taxas de falha padrão do componente e do sistema envolvidos, cujas taxas de falha são monitoradas constantemente. Consulte a seção de exclusão acima.

Se a Dell Technologies cancelar este Serviço de acordo com este parágrafo, ela enviará um aviso de cancelamento por escrito ao Cliente para o endereço indicado na fatura do Cliente. A notificação incluirá o motivo do cancelamento e a data de entrada em vigor do cancelamento, que não será em menos de 10 (dez) dias da data de envio do aviso de cancelamento pela Dell Technologies ao Cliente, a menos que alguma lei local exija outras provisões de cancelamento que não possam ser alteradas por contrato. Se a Dell Technologies cancelar este Serviço de acordo com este parágrafo, o Cliente não terá direito ao reembolso das tarifas pagas ou devidas à Dell Technologies.



i) Limitações Geográficas e Realocação.

- 1. A Dell Technologies pode não ser capaz de fornecer Serviços de Suporte a peças de reposição e no local em 4 horas com relação aos Equipamentos que estão fora da Área de Serviço da Dell Technologies. "Área da Dell Technologies Services" significa uma localização que está dentro de (i) um raio de 160 (cento e sessenta) quilômetros, ou 100 (cem) milhas, de carro de uma localização de serviço da Dell Technologies; e (ii) no mesmo país da localização de serviço da Dell Technologies, a menos que definido de outra forma no contrato vigente firmado com a Dell Technologies, hipótese em que a definição descrita no contrato em vigor prevalecerá. Para os Clientes situados na região EMEA, salvo indicação em contrário nesta Descrição de serviço ou no Contrato, o serviço no local está disponível até uma distância de 150 quilômetros da localização mais próxima de Logística da Dell Technologies (localização de retirada/entrega, ou PUDO). Antes da compra, entre em contato com seu representante de vendas para obter mais informações sobre a disponibilidade do serviço no local na região EMEA, de acordo com os locais de serviço da Dell Technologies.
- Este Serviço não está disponível em todos os locais. Se o Produto não estiver localizado na região geográfica que 2. corresponde ao local indicado nos registros de serviço da Dell Technologies para o Produto ou se os detalhes de configuração tiverem sido alterados e não comunicados à Dell Technologies, a Dell Technologies deverá primeiro requalificar seu Produto para o direito de suporte que você adquiriu antes de serem redefinidos os tempos de resposta aplicáveis para o Produto. As opções de Serviço, inclusive níveis de serviço, horas de suporte técnico e tempos de resposta no local, podem variar conforme a região geográfica e a configuração, e determinadas opções podem não estar disponíveis para compra na localização do Cliente. Portanto, entre em contato com seu representante de vendas para obter essas informações. A obrigação da Dell Technologies de prestação dos Serviços a Produtos realocados está sujeita a diversos fatores, incluindo, entre outros, disponibilidade local do Serviço, tarifas adicionais e inspeção e recertificação dos Produtos realocados, de acordo com as taxas atuais de consultoria de tempo e materiais da Dell Technologies. A menos que acordado de outro modo entre a Dell Technologies e o Cliente, nos casos em que as peças de serviço forem enviadas diretamente ao Cliente, o Cliente deverá ser capaz de receber as peças na localização de reparo dos Produtos. A Dell Technologies não será responsabilizada por atrasos no suporte devido à falha ou à recusa do Cliente em aceitar a remessa das peças. Os sistemas de armazenamento com vários componentes exigem contratos de opção de suporte ativos em todos os componentes de hardware e software do sistema a fim de receber todos os benefícios do contrato de suporte para a solução inteira. A menos que acordado de outra forma por escrito com o Cliente, a Dell Technologies reserva-se o direito de alterar o escopo dos Serviços de Suporte em sessenta (60) dias antes do prévio aviso por escrito ao Cliente.
- j) Ordem de Precedência. Salvo se acordado de outra forma por escrito entre as partes, se houver um conflito entre os termos de qualquer um dos documentos que compreendem este Contrato, os documentos prevalecerão na seguinte ordem: (i) esta Descrição de Serviço; (ii) o Contrato. As condições prevalecentes serão interpretadas na forma mais estrita possível para resolver o conflito preservando o maior número possível de condições não conflitantes, inclusive preservando as cláusulas não conflitantes no mesmo parágrafo, seção ou subseção.

ENTRE EM CONTATO CONOSCO

Para saber mais, entre em contato com o seu representante local ou revendedor autorizado.

Copyright © 2024 Dell Inc. ou suas subsidiárias. Todos os direitos reservados. Dell Technologies, Dell e outras marcas comerciais pertencem à Dell Inc. ou suas subsidiárias. Outras marcas comerciais podem ser marcas comerciais de seus respectivos proprietários. Publicado no Brasil.

A Dell Technologies acredita que as informações deste documento são precisas até a data da publicação. As informações estão sujeitas a alterações sem aviso prévio.



Dell Technologies Enterprise Systems Rail Sizing and Rack Compatibility Matrix

This document provides mounting features and key dimensions of the rack rails used for mounting many Dell Technologies enterprise systems and peripheral devices in a rack enclosure.



The information in this publication is provided "as is." Dell Inc. makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Use, copying, and distribution of any software described in this publication requires an applicable software license.

This document may contain certain words that are not consistent with Dell's current language guidelines. Dell plans to update the document over subsequent future releases to revise these words accordingly.

This document may contain language from third party content that is not under Dell's control and is not consistent with Dell's current guidelines for Dell's own content. When such third-party content is updated by the relevant third parties, this document will be revised accordingly.

Copyright © Jan 2025 | Version 5.0 Dell Inc. or its subsidiaries. All Rights Reserved. Dell Technologies, Dell, EMC, Dell EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be trademarks of their respective owners.

Contents

ntroduction	1
Considerations	1
Nounting interface	2
Rail types - System Installation Method	3
Cable Management Solutions	4
Backwards compatibility	5
elf-Adjusting Slide Feature	7
Definitions - Reference for Table 2	7

Figures

Figure 1.	Top view of right front EIA mounting flange	1
Figure 2.	System offset for round-hole racks	2
Figure 3.	Self-Adjusting Slide Feature	7

Tables

Table 1.	Dell Technologies server rails compatibility chart	. 5
Table 2.	DELL Technologies Rail Sizing Matrix	. 9
Table 3.	Dell Technologies rack compatibility matrix	22



Introduction

This document provides information about the mounting features and key dimensions of the rack rails used for mounting many Dell Technologies[™] enterprise systems and peripheral devices in a rack enclosure. This document also provides a compatibility summary for select Dell Technologies racks as well as some common third-party racks. Note that the product list is not all-inclusive and updates will be made as needed.

The dimensions provided in this document are for reference only. Some minor deviations due to manufacturing tolerances and variances should be expected.

Dell Technologies rail kits may not be compatible with racks from other vendors, however, all Dell Technologies rail kits are designed for compliance with all EIA-310-D and later revision specifications for 19-inch racks.

Considerations

Please pay attention to the footnotes indicated in the tables because they provide important information on using the rails in different racks and circumstances.

It is assumed that rack mount peripherals and cable bundles do not protrude into the space directly behind the systems.

Note that Dell Technologies rail kits with a Rail Identifier code have been designed to be compliant with the Server System Infrastructure (SSI) Specification for Computer Server Cabinet Enclosures & Racks, which specifies a minimum offset distance for return flanges on the rack mounting flanges to allow sufficient room for mounting the rail kits, as indicated in Figure 1. For more information about the Server System Infrastructure (SSI) Specification for Computer Server (SSI) Specification for Computer Server System Racks, see the SSI Forum at ssiforum.org.





Some third-party racks may not meet this requirement, and although Dell Technologies has made extensive efforts to accommodate as many third-party racks as possible, it is not feasible to provide a solution for every circumstance.



Rack Types - 2-post and 4-post

Dell Technologies rail kits install into two different rack types with various flange hole designs. These rack types are broken down in Table 2 into 4-post and 2-post styles. 4-post rack types contain vertical mounting flanges with either square-hole, unthreaded round-hole, or threaded round-hole designs as part of the rack and rail interface. 2-post rack types generally contain threaded round-hole designs and require users to mount the server in either the front or center mount orientations. Only stab-in static rail designs that support 2-post rack configurations may be installed into a 2-post rack and commonly require additional hardware to secure the rails to the rack. For more information, refer to the definitions section for Table 2.

Mounting interface

The ReadyRails^M II mounting interface supports tool-less installation in 4-post square-hole and unthreaded round-hole racks as well as native support for tooled installation in threaded-hole racks. Note that installing this mounting interface in a square-hole rack allows the bracket to be placed flush against the mounting post, while installation in a round-hole rack results in a slight offset of approx. 6 mm from the mounting post, which also results in an approx. 6 mm bezel offset; refer to Figure 2.



The original **ReadyRails** mounting interface is used for both static and sliding rails, and it supports tool-less installation in 4-post square-hole and unthreaded round-hole racks. Static ReadyRails kits also support tooled installation in threaded-hole racks and 2-post racks. When installed in unthreaded round-hole racks, the original ReadyRails will also have the 6 mm offset from the mounting post that was discussed in the previous ReadyRails II paragraph. In order to install sliding ReadyRails kits into a threaded-hole rack, adapter brackets are required. 1U and 2U adapter bracket kits are available that support systems ranging from 1U to 5U in height.

The adapter bracket kits include six brackets to accommodate different rail lengths, plus four sets of custom screws in 10-32, 12-24, M5 and M6 thread sizes. The design of the brackets has been optimized to limit the forward shift of the system in the rack to only 17.3 mm. Depending on the depth of the rack used and the position of the mounting rails within the rack, it may be necessary to remove the system's bezel in order to close the front door of the rack. For the front door to close with the system



bezel installed, a minimum clearance of 58 mm is needed between the back surface of the door panel and the front face of the EIA flange.

The **RapidRails**[™] mounting interface supports tool-less installation in 4-post square-hole racks only, while the **VersaRails**[™] mounting interface supports tooled installation in 4-post square-hole and unthreaded round-hole racks. Mounting the VersaRails in threaded-hole racks is not recommended and is not supported by Dell Technologies.

The Generic mounting interface encompasses all other mounting interfaces outside of the ones listed above. Unless indicated to be tool-less, tools are required for installation.

Rail types - System Installation Method

Drop-in/Stab-in rails (Combo Rail) are a feature rich rail solution that allows a system to be fully extended out of the rack for service and the user has the option to install the system into the rail using a drop-in method like the ReadyRails sliding rails, or a stab-in method like the ReadyRails static rails. Drop-in/Stab-in rails support CMA or SRB applications. CMA and SRB applications must be detached in order to remove the inner member from the rails.

A "drop-in" design means that the system is installed vertically into the rails by inserting the standoffs on the sides of the system into the "J-slots" in the inner rail members with the rails in the fully extended position. The recommended method of installation is to first insert the rear standoffs on the system into the rear J-slots on the rails to free up a hand and then rotate the system down into the remaining J-slots while using the free hand to hold the rail against the side of the system.

A "stab-in" design means that the inner (chassis) rail members must first be attached to the sides of the system and then inserted into the outer (cabinet) members installed in the rack. For systems that are 2U and larger, it is recommended that two people perform this operation.

Sliding rails allow a system to be fully extended out of the rack for service. Most sliding rails support Cable Management Arms (CMAs) which enable the system to be extended out of the rack without disconnecting data/power cables at the rear of the system.

Unless otherwise indicated, all sliding rails are drop-in sliding rail design.

Static rails typically do not support the ability to service the system in the rack and are not compatible with the CMA. However, they do offer more flexibility in the types of racks and installations supported. Generally, there are two types of static rails: stab-in static and L-bracket static.

Stab-in static rails require the inner (chassis) rail members must first be attached to the sides of the system and then inserted into the outer (cabinet) members installed in the rack. For systems that are 2U and larger, it is recommended that two people perform this operation.

Stab-in Sliding rails require the inner (chassis) rail members must first be attached to the sides of the system and then inserted into the outer (cabinet) members installed in the rack and allow a system to be portion extended out of the rack for service. For systems that are 2U and larger, it is recommended that two people perform this operation. Most stab-in Sliding rails are compatible with CMA and SRB solutions.

L-bracket static rails do not support the ability to fully extend a system out of the rack into a service position. These rails typically are not compatible with cable management solutions unless otherwise indicated. Typically, equipment supported by L-bracket are customer serviceable from the front or rear of the rack.



Cable Management Solutions

To help manage the numerous cables associated with rack-mounted servers, a Cable Management Arm (CMA) or Strain Relief Bar (SRB) can be used. An optional CMA is offered with most sliding rails. CMAs attach on either the right or left side without tools.

Cable management arm (CMA) is a cable management accessory which connects to the rails behind the system. It allows a fully cabled system to be extended out of the rack into a service position.

Strain relief bar (SRB) is a cable management solution, which in most cases, attaches to the back of the rails via the strain relief bar brackets. Cables from the back of the chassis are placed across the top of the SRB and secured by straps.

SRBs are offered for select systems as an optional method for managing cables at the rear of the system due to the potential of a cable bundle size that exceeds the capacity of the CMA. The rail depth with a SRB is significantly less than that of a CMA, which in many cases, enables fitment of the rails in shallow racks. Cable service loops are required for systems on sliding rails to fully extend out of the rack for service.

Note that using a CMA or SRB with a deeper system may interfere with access to power distribution units (PDUs) in certain racks. If a configuration does not require CMA support, then the outer CMA mounting brackets can be removed from some of the sliding rail kits to reduce the overall length of the rails and eliminate potential interference with rear-mounted PDUs or the rack rear door.



Backwards compatibility

Some systems may offer backward compatibility with the rail kits from previous-generation systems. This is not always possible, because changes to chassis features, dimensions or weight can prevent older rail kits from being used with newer systems. Please refer to Table 1 for cross-generational compatibility of Dell Technologies servers and rails.

17 th Generation	Backwards compatibility with 16 th generation rails/CMAs			
product	Sliding rails	СМА	Static rails	
R670	✓	✓	1	
R770	✓	✓	1	
R570	✓	✓	1	
R470	X	X	X	
R7715	✓	✓	1	
R7725	✓	✓	1	
R6715	✓	✓	✓	
R6725	✓	✓	✓	

Table 1. Dell Technologies server rails compatibility chart

16 th Generation	Backwards compatibility with 15 th generation rails/CMAs			
product	Sliding rails	СМА	Static rails	
R260	N/A	N/A	✓	
R360	1	✓	✓	
R760	X	√*	X	
R660	✓	√*	✓	
R7615	X	X	X	
R7625	✓	√*	✓	
R6615	X	X	X	
R6625	✓	√*	✓	
R760xd2	X	X	X	
R660xs	✓	√*	✓	
R760xs	✓	√*	✓	
R760xa	X	N/A	X	
R860	N/A	N/A	N/A	
R960	N/A	N/A	N/A	



15 th Generation	Backwards compa	tibility with 14th gen	eration rails/CMAs
product	Sliding rails	СМА	Static rails
R250	N/A	N/A	✓
R350	✓	1	✓
R650xs	✓	✓	✓
R450/R450xs	✓	✓	✓
R750xs	X	✓	X
R550/R550xs	X	✓	X
R6515	✓	✓	✓
R6525	N/A	N/A	N/A
R650	X	✓	X
R7515	✓	✓	✓
R7525	X	X	X
R750	✓	1	✓
R750xa	N/A	N/A	N/A

14 th Generation	Backwards com	npatibility with 13 th gene	eration rails/CMAs
product	Sliding rails	СМА	Static rails
R240	N/A	N/A	✓
R340	X	✓	✓
R440	X	✓	✓
R540/R540xd	✓	✓	✓
R640	✓	✓	✓
R740/R740xd	✓	✓	✓
R740xd2	N/A	N/A	N/A
R840/940xa	X	X	X
R940	X	✓	N/A
C4140	N/A	N/A	✓
C64xx	N/A	N/A	✓
T440	Х	✓	N/A
T640	\checkmark	\checkmark	N/A

✓ - Compatible
 X - Not compatible
 *Only with the previous generation sliding rail

6 D≪LLTechnologies



Self-Adjusting Slide Feature

For many 1U and 2U systems, rails have been standardized with a slim design that holds a wide system chassis to accommodate more features and functions. They also have a self-adjusting slide feature that accommodates different depths of systems, offering compatibility across multiple platform models. Refer to Figure 3 for an illustration of how the self-adjusting slide feature works.



The rail adjustability range when the rails are installed in a rack is the same regardless of system depth since the feature is not utilized until a system is installed. If the system being installed in the rails requires this feature, the minimum rail adjustability limit is increased by the amount of travel the slide body needs to slide back to support the system. The minimum rail adjustability limit is documented in the resources listed at the end of this notice.

Users who have systems that utilize the feature might observe a slight amount of additional resistance from a spring in each rail when the system is almost completely installed in the rack. For most rails, the instance when the resistance is observed is within the final 55 mm of translation before the slam latch is engaged with the rail.

The rail slide-adjusting feature can be found on both sliding and drop-in/stab-in rail types. The rail adjustability range (mm) values listed in Table 2 for products that utilize this rail feature have been flagged with a footnote.

Definitions - Reference for Table 2

Rail identifier is a two-character code used on most rail kits to indicate compatibility between rails and systems. The twocharacter code consists of a letter followed by a one or two-digit number. It is typically located on a front inside surface on both the left and right sliding rail and drop-in/stab-in rail members. If there is a component of the rail kit that is attached to the chassis prior to installing the system into a rack, such as with the stab-in static rails, the identifier is located closer to the center of the component.

Square-hole describes a 4-post rack mounting flange type where rails utilize Square holes sized according to EIA-310-D standard for mounting.

Round-hole describes a 4-post rack mounting flange type where rails utilize Unthreaded-Round holes sized according to EIA-310-D standard for mounting.

7



Threaded-hole describes a 4-post rack mounting flange type where rails utilize Threaded-round holes for mounting. Threaded-round holes may require additional hardware for mounting and hardware may vary by thread type. See footnotes in table 2 for specific information on threaded-round hole mounting.

Mounting interface describes the type of rail bracket design used for mounting the rail in the rack.

Rail adjustability range represents the allowable distance between the outside-facing surfaces of the front and rear mounting posts of the rack when a system is fully installed. This does not include the portion of the rail kit or other rail components that may extend beyond the mounting posts.

Rail depth represents the minimum depth of the rail as measured from the rack front mounting posts when the rail rear bracket is positioned all the way forward. The rail may extend beyond the rear bracket, particularly for sliding rail kits to support CMA or SRB attachment. In some instances, the chassis may extend beyond the minimum rail depth, and in such cases, please refer to the footnotes in Table 2.

Table 2. DELL Technologies Rail Sizing Matrix

							Rack t	ypes sup	ported			Rail a	djustabi	lity rang	e (mm)		Rail dep	th (mm)
		Product	Rail identifier	Mounting interface	Rail type		4-Post		2-	Post	Squ	are	Ro	und	Thre	aded	without	with
						Square	Round	Thread	Flush	Center	Min	Max	Min	Max	Min	Max	CMA/SRB	CMA(SRB)
			Α7	ReadyRails II	Sliding	~	~	√ a,c,d	X	x	631	868	617	861	631	883	720 ^b	845
		R640 (8-HDD)	A8	ReadyRails	Stab-in Static	~	~	√ a,c	√a,c	√c	608	879	594	872	610	898	622 ⁿ	-
			A10	Generic Tool-less	Drop-in/ Stab-in	~	~	✓ ₩	х	x	559	931	559	931	559	931	720 ^b	845
			Α7	ReadyRails II	Sliding	~	~	√ a,c,d	х	x	681 ^p	868	667 ^p	861	681 ^p	883	770 ^b	895
		R640 (4-HDD/10-HDD)	A8	ReadyRails	Stab-in Static	~	~	√ a,c	√ a,c	√ c	608	879	594	872	610	898	622 ⁿ	-
			A10	Generic Tool-less	Drop-in/ Stab-in	~	~	✓ ₩	х	x	613 ^p	931	613 ^p	931	613 ^p	931	770 ^b	895
/RS	:dge TM	R6525 (8-HDD)	A15	ReadyRails II	Sliding	~	~	√ a,c,d	х	x	631	868	617	862	631	884	736 ^b	862 (770/792)
SER	owerE	R650 (8-HDD) R660/R6615/R6625 (0-HDD/8-HDD)	A14	ReadyRails	Stab-in Static	~	×	√ a,c	√a,c	√ c	608	880	594	870	605	893	622 ⁿ	-
	đ		A16	Generic Tool-less	Drop-in/ Stab-in	~	~	✓ ₩	x	x	559	994	559	994	559	944	736 ^b	862 (770/792)
		R6525 (4-HDD/10-HDD)	A15	ReadyRails II	Sliding	~	×	√ a,c,d	x	x	631	868	617	862	631	884	787 ^ь	913 (821/843)
		R650 (4-HDD/10-HDD) R660/ R6615/R6625	A14	ReadyRails	Stab-in Static	~	*	√ a,c	√a,c	√c	608	880	594	870	605	893	622 ⁿ	-
		(4-000/10-000/23/L0-2) R670/R6725/R6715/R470	A16	Generic Tool-less	Drop-in/ Stab-in	~	~	1	x	X	610	994	610	994	610	994	787 ^ь	913 (821/843)
		R340 (8-HDD)	A12	ReadyRails II	Sliding	~	~	√ a,c,d	X	x	631	868	617	861	631	883	720 ^b	845
		R360 (8-HDD)	A8	ReadyRails	Stab-in Static	~	✓	√ a,c	√ a,c	√ c	608	879	594	872	610	898	622 ⁿ	-



					Rack t	ypes sup	ported			Rail a	djustabi	lity rang	e (mm)		Rail dep	th (mm)
Product	Rail identifier	Mounting interface	Rail type		4-Post		2-1	Post	Squ	are	Ro	und	Thre	aded	without	with
				Square	Round	Thread	Flush	Center	Min	Max	Min	Max	Min	Max	CMA/SRB	CMA(SRB)
R340 (4-HDD)	A12	ReadyRails II	Sliding	~	×	√ a,c,d	X	X	681 ^p	868	667 ^p	861	681 ₽	883	770 ^b	895
R360 (4-HDD)	A8	ReadyRails	Stab-in Static	~	~	√ a,c	√ a,c	√c	608	879	594	872	610	898	622 ⁿ	-
R440 (8-HDD) R450 (8-HDD) R6415 (8-HDD)	A8	ReadyRails	Stab-in Static	~	~	√ a,c	√ a,c	√ c	608	879	594	872	610	898	622 ⁿ	-
R6515 (8-HDD) R650xs (0-HDD/8-HDD) R660xs(0-HDD/8-HDD)	A11	Generic Tool-less	Drop-in/ Stab-in	~	~	✓ ₩	x	x	559	931	559	931	559	931	720 ^b	845 (761/783)
R440 (4-HDD/10-HDD) R450 (4-HDD) R6415 (4-HDD/10-HDD)	A8	ReadyRails	Stab-in Static	~	~	√a,c	√ a,c	√c	608	879	594	872	610	898	622 ⁿ	-
R6515 (4-HDD/10-HDD) R650xs (4-HDD/10-HDD) R650xs (8-HDD NVME) R660xs (4-HDD/10-HDD)	A11	Generic Tool-less	Drop-in/ Stab-in	*	*	*	X	X	609 ^p	931	609 ^p	931	609 ^p	931	770 ^b	895 (811/833)
	B6	ReadyRails II	Sliding	~	✓	√ a,c,d	X	x	631	868	617	861	631	883	714 ^b	845
R540/R540xd R740/R740xd/ P7415/P7425/P7515	B4	ReadyRails	Stab-in Static	~	✓	√ a,c	√ a,c	√c	608	879	594	872	610	898	622 ⁿ	-
K/415/K/425/K/515	B13	Generic Tool-less	Drop-in/ Stab-in	~	×	✓ *	X	x	559	931	559	931	559	931	714 ^b	845
	B20	ReadyRails	Stab-in Static	~	~	√ a,c	√ a,c	√ c	608	879	594	872	610	898	622 ⁿ	-
R550 R750xs	B21	ReadyRails II	Sliding	~	~	√ a,c,d	X	X	631	868	617	861	631	883	714 ^b	845 (749/769)
K760XS	B22	Generic Tool-less	Drop-in/ Stab-in	~	~	~	x	x	559	931	559	931	559	931	714 ^b	845 (749/769)
87565	B6	ReadyRails II	Sliding	~	~	√ a,c,d	X	X	685 ^p	868	671 ^p	861	685 ^p	883	766 ^b	895 (802/822)
R7525 R750	B4	ReadyRails	Stab-in Static	~	~	√ a,c	√ a,c	√c	608	879	594	872	608	898	622 ⁿ	-
	B13	Generic Tool-less	Drop-in/ Stab-in	~	~	√ ₩	X	X	609 ^p	931	609 ^p	931	609 ^p	931	779 ^ь	899 (802/822)

10

D&LLTechnologies





11

OTOCO

Fls.<u>1018</u> Mov. <u>33</u>

ADO DO

D&LLTechnologies

	Dell Technologies Ente	erprise Syst	ems Rail Sizin	g and Rack	Compa	tibility	Matrix										
						Rack t	ypes sup	ported			Rail ac	ljustabi	lity rang	e (mm)		Rail dep	oth (mm)
	Product	Rail identifier	Mounting	Rail type		4-Post		2-	Post	Squ	are	Ro	und	Thre	aded	without	with
					Square	Round	Thread	Flush	Center	Min	Max	Min	Max	Min	Max	CMA/SRB	CMA(SRB)
	FV2/FV2-	B10	ReadyRails II	Sliding	~	✓	√a,c,d	Х	х	677	815	665	809	677	830	836	888
	FXZ/FXZS	B11	ReadyRails II	Stab-in Static	~	1	√ a,c	Х	x	644	916	632	910	644	930	828	-
	C4130/C4140	А9	ReadyRails II	Stab-in Static ^h	~	~	√a,c,d	X	X	643	916	631	910	643	930	766	-
	T640	C4	ReadyRails II	Sliding	~	~	√ a,c,d	X	x	686	756	672	749	686	771	756	840
	T440	C2	ReadyRails II	Sliding	~	×	√a,c,d	X	x	686	756	672	749	686	771	760	840
	VRTX	С3	ReadyRails II	Sliding	~	✓	√a,c,d	X	x	608	915	594	908	608	930	756	845
	R240/R250/R260	A4	ReadyRails	Stab-in Static	~	~	√a,c	√ a,c	√c	608	879	594	872	610	898	622 ⁿ	-
	M1000a	-	RapidRails	L-Bracket Static	~	X	x	X	x	712	755	-	-	-	-	703	-
	MICODE	-	VersaRails	L-Bracket Static	~	~	x	X	x	706	755	706	755	-	-	703	-
	MX7000	C5	ReadyRails II	L-Bracket Static	~	~	x	X	x	592	876	578	869	-	-	m	(901)
		A12	ReadyRails II	Sliding	~	~	√ a,c,d	X	x	681 ^p	868	667 ^p	861	681 ^p	883	770 ^b	895
	XR2	A4	ReadyRails	Stab-in Static	~	✓	√ a,c	√ a,c	✓c	608	879	594	872	610	898	622 ⁿ	-
		-	Generic	Stab-in Static ^t	√g	√g	√g	X	x	464	766	464	766	464	766	464 ⁿ	-
	VD44/VD42	A20	Generic Tool-less	Stab-in Sliding	~	~	~	√a	√a	472	757	472	757	472	757	445 ^b	605 (498/520)
	ΔΚΤΙ/ΔΚΊΖ	A21 ^s	Generic Tool-less	Stab-in Sliding	X	~	x	X	x	-	-	458	589	-	-	-	-
U e	C1100	-	Generic Tool-less	Sliding	~	1	x	X	x	665	950	665	950	-	-	685	-
erEdg	C2100	-	Generic	Sliding	~	1	~	X	x	664	1110	664	1110	664	1110	720	-
Pow	C410x	-	VersaRails	Stab-in Static	~	~	x	X	x	737	972	737	972	-	-	734	-

D&LLTechnologies

12

OTOC

Fls.1019 Mov. 33

	Ĭ					Rack t	ypes sup	ported			Rail ac	djustabi	lity rang	e (mm)		Rail dep	th (mm)
	Product	Rail identifier	Mounting interface	Rail type		4-Post		2-	Post	Squ	are	Ro	und	Thre	aded	without	with
					Square	Round	Thread	Flush	Center	Min	Max	Min	Max	Min	Max	CMA/SRB	CMA(SRB)
	С5ххх	-	Generic Tool-less	L-Bracket Static	×	✓	X	X	X	708	947	708	947	-	-	705	-
	C64xx C65xx C66xx	-	Generic Tool-less	L-Bracket Static	~	~	X	X	X	609ª	917	609ª	917	-	-	-	-
	C8000	-	Generic Tool-less	L-Bracket Static	~	~	X	x	X	708	946	708	946	-	-	713	-
		B20	ReadyRails	Stab-in Static	~	1	√ a,c	√ a,c	√c	608	879	594	872	610	898	622 ⁿ	-
	HS5620	B21	ReadyRails II	Sliding	~	~	√ a,c,d	x	X	631	868	617	861	631	883	714 ^b	845 (749/769)
		B22	Generic Tool-less	Drop-in/ Stab-in	~	~	~	X	X	559	931	559	931	559	931	714 ^b	845 (749/769)
P.	HS5610 Cold Aisle	A22	ReadyRails II	Stab-in, Static	~	~	√ a,c,d	x	x	608	916	594	909	610	641	897	-
Edge C		A8	ReadyRails	Stab-in Static	~	~	√ a,c	√ a,c	√c	608	879	594	872	610	898	622 ⁿ	-
Power	133010(0-000/8-000)	A11	Generic Tool-less	Drop-in/ Stab-in	~	~	√ ₩	x	X	559	931	559	931	559	931	720 ^b	845 (761/783)
		Α8	ReadyRails	Stab-in Static	~	~	√ a,c	√ a,c	√c	608	879	594	872	610	898	622 ⁿ	-
	(עעחייטו עעמייא)	A11	Generic Tool-less	Drop-in/ Stab-in	~	~	√ ₩	x	x	609 ^p	931	609 ^p	931	609 ^p	931	770 ^b	895 (811/833)
XR		A23			✓	✓	√a	X	x	552	763	552	763	552	763	559 ^b	
erEdge	XR4000R	A24	Generic Tool-less	Stab-in Sliding	~	~	√a	X	X	342	554	342	554	342	554	358 ^b	-
Powe		A25 ^s			x	~	x	x	x	-	-	426	569	-	-	-	

13

D&LLTechnologies





14

OTOCO

Fls.<u>1021</u> Mov. <u>33</u>

ADO DO

D&LLTechnologies

							Rack t	ypes sup	ported			Rail ac	djustabi	lity rang	e (mm)		Rail dep	th (mm)
		Product	Rail identifier	Mounting interface	Rail type		4-Post		2-	Post	Squ	are	Ro	und	Thre	aded	without	with
						Square	Round	Thread	Flush	Center	Min	Max	Min	Max	Min	Max	CMA/SRB	CMA(SRB)
		1081AD/2161AD 1082DS/2162DS 4322DS	A5	ReadyRails	Stab-in Static	~	~	✓	~	✓	496	770	482	763	488	794	506 ^Q	-
	KVM	180AS/2160AS 2161DS/2161DS-2 4161DS	-	Generic	Stab-in Static	*	~	~	~	x	686	737	686	737	686	737	686	-
		2321DS	-	Generic	Stab-in Static	~	*	~	~	X	533	737	533	737	533	737	533	-
		PC8132/PC8132F PC8164/PC8164F	Α5	ReadyRails	Stab-in Static	×	×	×	✓	×	496	770	482	763	488	794	506 ^Q	-
		S4820T/S6000	A5	ReadyRails	Stab-in Static	×	×	✓	~	×	496	770	482	763	488	794	506 ^Q	-
		\$5000	-	Generic	Stab-in Static	~	~	~	X	X	680	830	680	830	680	830	680	-
HES		Z9100	A5	ReadyRail	Stab-in Static	~	~	~	✓	×	496	770	482	763	488	794	506 ^Q	-
SWITC		S4248	A5	ReadyRail	Stab-in Static	~	~	~	~	~	496	770	482	763	488	794	506 ^Q	-
	ing	S41xx	A5	ReadyRail	Stab-in Static	~	~	~	~	×	496	770	482	763	488	764	506 ^Q	-
	twork	S4048/S4048T	A5	ReadyRail	Stab-in Static	~	✓	~	~	1	496	770	482	763	488	764	506 ^Q	-
	Ne	S6010	A5	ReadyRail	Stab-in Static	~	✓	~	~	~	496	770	482	763	488	764	506 ^Q	-
		S3048	A5	ReadyRail	Stab-in Static	~	✓	1	~	~	496	770	482	763	488	764	506 ^Q	-
		S6100	B9	ReadyRails II	L-Bracket Static	~	✓	√ a,c,d	X	X	595	914	581	907	595	929	600	-
		S6100NEBS	-	Generic	Stab-in Static	X	X	x	~	x	-	-	-	-	-	-	-	-
		N2128PX-ON	-	Generic	Stab-in Static	X	X	X	✓	X	-	-	-	-	-	-	-	-
		N3132PX-ON	Α5	ReadyRails	Stab-in Static	✓	✓	~	×	✓	496	770	482	763	488	764	506 ^Q	-

OTOC

Fls.1022 Mov. <u>33</u>

D&LLTechnologies

Inserido ao protocolo 22.951.206-4 por: Pedro Henrique Golin Linhares em: 31/03/2025 12:49. A autenticidade deste documento pode ser validada no endereço: https://www.eprotocolo.pr.gov.br/spiweb/validarDocumento com o código: fc1c44dfe0686d5547ffd691199be699.

15

						Rack t	ypes sup	ported			Rail a	djustabi	lity rang	e (mm)		Rail dep	th (mm)
	Product	Rail identifier	Mounting interface	Rail type		4-Post		2-	Post	Squ	are	Ro	und	Thre	aded	without	with
					Square	Round	Thread	Flush	Center	Min	Max	Min	Max	Min	Max	CMA/SRB	CMA(SRB)
	N1108T/N1108P	-	Generic	Stab-in Static	x	X	x	✓	x	-	-	-	-	-	-	-	-
	N1124T/N1124P	-	Generic	Stab-in Static	X	X	X	~	X	-	-	-	-	-	-	-	-
	N1148T/N1148P	-	Generic	Stab-in Static	X	X	X	~	X	-	-	-	-	-	-	-	-
	N3024/N3048	A5	ReadyRails	Stab-in Static	~	~	1	~	~	496	770	482	763	488	764	506 ^Q	-
	S5148	A5	ReadyRails	Stab-in Static	~	×	~	×	×	496	770	482	763	488	764	506 ^Q	-
	S31xx	A5	ReadyRail	Stab-in Static	*	×	*	×	×	496	770	482	763	488	764	506 ^Q	-
	N30xx	A5	ReadyRail	Stab-in Static	~	~	1	~	~	496	770	482	763	488	764	506 ^Q	-
	P7010	B6	ReadyRails II	Sliding	*	~	√ a,c,d	X	X	631	868	617	861	631	883	714 ^b	845
	K7910	B4	ReadyRails	Stab-in Static	~	×	√ a,c	√ a,c	✓c	608	879	594	872	610	898	622	-
	Precision 3930 Rack	A4	ReadyRails	Stab-in Static	*	~	√ a,c	√ a,c	√c	608	879	594	872	610	898	622 ⁿ	-
10		B6	ReadyRails II	Sliding	*	×	√ a,c,d	X	X	631	868	617	861	631	883	714 ^b	845
	Precision 7920 Rack	B4	ReadyRails	Stab-in Static	*	*	√ a,c	√ a,c	√c	608	879	594	872	610	898	622 ⁿ	-
RKST⊿		B13	Generic Tool-less	Drop-in/ Stab-in	*	×	✓ ₩	X	X	607 ^p	931	607 ^p	931	607 ^p	931	714 ^b	845
MO		B20	ReadyRails	Stab-in Static	~	~	√ a,c	√ a,c	✓c	608	879	594	872	610	898	622 ⁿ	-
	Procision 7960 Pack	B21	ReadyRails II	Sliding	~	✓	√ a,c,d	X	X	685 ^p	868	671 ^p	861	685 ^p	883	766 ^b	895 (802/822)
	FIECISION / 900 RdCK	B22	Generic Tool-less	Drop-in/ Stab-in	~	~	~	x	x	609 ^p	931	609 ^p	931	609 ^p	931	766 ^b	895 (802/822)
	T7600/T7610	C2	ReadyRails II	Sliding	~	~	√ a,c,d	X	X	686	756	672	749	686	771	760	840

16

OTOC

Fls.<u>1023</u> Mov. <u>33</u>

D&LLTechnologies

						Compa	Rack t	ypes sup	ported			Rail ac	ljustabi	lity rang	e (mm)		Rail dep	th (mm)
		Product	Rail identifier	Mounting interface	Rail type		4-Post		2-	Post	Squ	are	Ro	und	Thre	aded	without	with
						Square	Round	Thread	Flush	Center	Min	Max	Min	Max	Min	Max	CMA/SRB	CMA(SRB)
		R5500/R7610	B2	ReadyRails	Sliding	✓	×	√f	x	x	686	883	672	876	651	897	755 ^b	883
		FPM185 (without KVM)	-	ReadyRails II	Sliding	~	×	√ a,c,d	X	x	604	900	590	893	604	914	-	611
į	Ŵ	FPM185 (with KVM)	-	ReadyRails II	Sliding	✓	×	√ a,c,d	X	x	705	900	691	893	705	914	-	715
:	2	17FP	-	RapidRails	Sliding	~	X	X	X	X	714	755	-	-	-	-	-	787
		.,,,,	-	VersaRails	Sliding	~	~	X	x	x	709	755	709	755	-	-	-	787
	CHU	Dell Rack Mount UPS Family	В5	ReadyRails	Stab-in Static	~	×	√f	X	x	518	769	504	762	483	783	526	-
	ОТНЕК	1U Fixed Equipment Shelf	Α4	ReadyRails	Stab-in Static	>	*	√ a,c	√ a,c	√c	608	879	594	872	610	898	622	-
	PowerMax	PROXPE1 PROXPE2	B14	-	L-Bracket Static	√u	✓u	√u,v	x	x	558	914	558	914	558	914	600	1015
		NIX 2 200 (NIX 400	Α7	ReadyRails II	Sliding	~	×	√ a,c,d	x	x	631	868	617	861	631	883	720 ^b	845
		NX3300/NX400	A8	ReadyRails	Stab-in Static	~	×	√ a,c	√ a,c	✓c	608	879	594	872	610	898	622	-
RAGE	ž	NY3200	B6	ReadyRails II	Sliding	~	×	√ a,c,d	X	X	631	868	617	861	631	883	714 ^b	845
ѕто	Vault	14,5200	B4	ReadyRails	Stab-in Static	✓	×	√ a,c	√ a,c	✓c	608	879	594	872	610	898	622	-
	ower	NY3500 Controller	Α3	ReadyRails	Sliding	×	×	√e	X	x	686	883	672	876	651	897	714 ^b	835
			Α4	ReadyRails	Stab-in Static	~	 Image: A start of the start of	√ a,c	√ a,c	√c	608	879	594	872	610	898	622	-
		NX3500 UPS	Α4	ReadyRails	Stab-in Static	~	×	√ a,c	√ a,c	√c	608	879	594	872	610	898	622	-
		DX6000G	Α4	ReadyRails	Stab-in Static	~	~	√ a,c	√ a,c	√c	608	879	594	872	610	898	622	-

17

D&LLTechnologies



					Rack t	ypes sup	ported			Rail a	djustabi	lity rang	e (mm)		Rail dep	th (mm)
Product	Rail identifier	Mounting interface	Rail type		4-Post		2-	Post	Squ	are	Ro	und	Thre	aded	without	with
				Square	Round	Thread	Flush	Center	Min	Max	Min	Max	Min	Max	CMA/SRB	CMA(SRB)
	A6	ReadyRails	Stab-in Static	~	~	√ a,c	√ a,c	√c	508 ^c	751	494 ^c	744	519 ^c	762	515 ^c 376 ^d	-
NY 300/DY 600/S	Α3	ReadyRails	Sliding	×	✓	√e	X	x	686	883	672	876	651	897	714 ^b	835
14300/0700043	Α4	ReadyRails	Stab-in Static	×	✓	√a,c	√ a,c	√c	608	879	594	872	610	898	622	-
NX 3000/DX 6000	B1	ReadyRails	Sliding	~	✓	√f	X	x	692	756	678	749	657	770	751	840
NX3000/DX0000	A2	ReadyRails	Stab-in Static	*	~	√ a,c	√ a,c	√c	588	828	574	821	592	846	608	-
NX3100/DL2200	В3	ReadyRails	Sliding	*	~	√f	X	X	686	883	672	876	651	897	714 ^b	845
DX6012S/DR4000	B4	ReadyRails	Stab-in Static	*	~	√ a,c	√ a,c	√c	608	879	594	872	610	898	622	-
MD3060e/MD3660	-	VersaRail	L-Bracket Static	*	~	x	X	X	611	791	611	791	-	-	620	-
	В9	ReadyRails II	L-Bracket Static	*	×	√ a,c,d	X	X	595	914	581	907	595	929	600	-
MD12xx/14xx/32xx/ 36xx/NX36xx	-	RapidRails	L-Bracket Static	~	X	x	X	x	732	758	-	-	-	-	729	-
	-	VersaRails	L-Bracket Static	*	×	x	X	X	714	758	714	758	-	-	721	-
MD1120	-	RapidRails	L-Bracket Static	*	X	x	X	X	732	759	-	-	-	-	729	-
MDTT20	-	VersaRails	L-Bracket Static	~	×	x	X	x	714	759	714	759	-	-	721	-
WD1000/WD2000	-	RapidRails	L-Bracket Static	~	X	x	X	x	732	758	-	-	-	-	735	-
MD 1000/MD 3000	-	VersaRails	L-Bracket Static	~	~	X	X	X	714	758	714	758	-	-	735	-
	Β7	ReadyRails	Stab-in Static	~	~	√ a,c	√ a,c	√c	588	828	574	821	592	846	608	-
	-	RapidRails	Sliding	~	X	X	X	X	722	750	-	-	-	-	792	870
F V I I 4 I / F V I I 4A	-	VersaRails	Sliding	✓	~	x	X	X	701	745	701	745	-	-	792	870

D&LLTechnologies



						Rack t	ypes sup	ported			Rail a	djustabi	lity rang	e (mm)		Rail dep	th (mm)
	Product	Rail identifier	Mounting interface	Rail type		4-Post		2-	Post	Squ	are	Ro	und	Thre	aded	without	with
					Square	Round	Thread	Flush	Center	Min	Max	Min	Max	Min	Max	CMA/SRB	CMA(SRB)
		-	RapidRails	L-Bracket Static	~	x	x	X	x	729	755	-	-	-	-	732	-
	PV1241	-	VersaRails	L-Bracket Static	~	~	х	X	х	711	755	711	755	-	-	732	-
		A1	ReadyRails	Sliding	~	~	√e	X	х	692	756	678	749	657	770	768 ^b	887
	FS7500 Controller	A2	ReadyRails	Stab-in Static	~	~	√a,c	√ a,c	√c	588	828	574	821	592	846	608	-
	FS7500 UPS	A4	ReadyRails	Stab-in Static	~	~	√ a,c	√ a,c	√c	608	879	594	872	610	898	622	-
.ogic™		B9	ReadyRails II	L-Bracket Static	~	~	√a,c,d	X	x	595	914	581	907	595	929	600	-
qualL	FS76xx/PS41xx PS61xx	-	RapidRails	L-Bracket Static	~	х	x	X	x	732	758	-	-	-	-	729	-
		-	VersaRails	L-Bracket Static	~	~	x	X	x	714	758	714	758	-	-	721	-
	PS6500/6510	-	ReadyRails	Sliding	~	~	√a,c	X	x	597	793	583	786	605	817	885	885
	PS4000/6000/6010	-	Generic	L-Bracket Static	~	√a	√a	X	x	616	914	616	914	616	914	616	-
	558000	B6	ReadyRails II	Sliding	~	~	√ a,c,d	X	x	631	868	617	861	631	883	714 ^b	845
	508000	B4	ReadyRails	Stab-in Static	~	~	√a,c	√ a,c	√c	608	879	594	872	610	898	622	-
ent™	SC20xx/SC40xx	-	Generic	L-Bracket Static	~	~	✓A	X	x	611	914	614	914	614	914	-	-
mpell		B9	ReadyRails II	L-Bracket Static	~	~	√a,c,d	X	x	595	914	581	907	595	929	600	-
ell Co	SC2xx/FS86xx	-	RapidRails	L-Bracket Static	~	X	x	X	x	732	758	-	-	-	-	729	-
Δ		-	VersaRails	L-Bracket Static	~	~	x	x	x	714	758	714	758	-	-	721	-
	SCV30xx SC50xx SC7020	В9	ReadyRails II	L-Bracket Static	~	~	√ a,c,d	x	x	595	914	581	907	595	929	600	-

19

OTOC

Fls.1026 Mov. 33

D&LLTechnologies


_	Dell Technologies Ente	erprise Syst	ems Rail Sizin	g and Rack	Compa	tibility	Matrix										
		D. 11				Rack t	ypes sup	ported			Rail a	djustabi	ility rang	e (mm)		Rail dep	th (mm)
	Product	Rail identifier	Mounting	Rail type		4-Post		2-	Post	Squ	are	Ro	und	Thre	aded	without	with
					Square	Round	Thread	Flush	Center	Min	Max	Min	Max	Min	Max	CMA/SRB	CMA(SRB)
	Series 40	-	Generic	Sliding	~	√g	√g	X	x	669	923	669	923	707 ^g	961 ^g	693	-
	Fibre Channel	-	Generic	Stab-in Static ^h	~	×	*	X	X	606	910	606	910	606	910	598	-
	SAS (new rails)	-	Generic	Stab-in Static ^h	~	×	X	X	X	606	910	606	910	606	910	598	-
	SAS (old rails)	-	Generic	Stab-in Static ^h	~	~	*	X	x	682	885	682	885	682	885	598	-
	NAS Gen3	-	Generic	Sliding	√i	√i	√i	X	X	652	854	652	854	652	854	810	-
				1	1						l	<u> </u>	l	1		1	

D&LLTechnologies

Notes:

- ^a Minor conversion required
- ^b With CMA brackets removed
- ^c Mounting screws not included in the kit
- ^d Mounting screw head diameter must be 10 mm or less
- ^e Requires the 1U Threaded Rack Adapter Brackets Kit (Dell PN 8Y19G), which shifts the system forward in the rack by 17.3 mm
- ^f Requires the 2U Threaded Rack Adapter Brackets Kit (Dell PN PKCR1), which shifts the system forward in the rack by 17.3 mm
- ^g Requires adapter kit (included)
- ^h System is serviceable while in the rack
- ⁱ Requires additional rail guide (included in kit) for full serviceability
- ^j With middle brackets removed
- ^k With rear brackets removed (applies to 2-post or cantilever mount only)
- ¹ SRB is staged furthest to the rack door
- ^m Rail depth is dependent on spacing between the front and rear mounting flanges of the rack Add amount below based on flange type:
 - Square hole (5.7mm)
 - Round hole (11.8mm)
- ⁿ Rail depth represents cabinet assembly only and does not represent inner rail component that attaches to chassis
- ° Footnote intentionally left blank
- ^p Chassis type utilizes the Self-Adjusting Rail Feature to install properly into rack
- ^q Depth maybe greater based on rail adjustability range
- $^{\rm r}$ Rail threaded mount only compatible with #10-32 thread type
- ^s Rail is only intended for use with ruggedized transit case (Pelican custom rack 25-036329-01)
- ^t Rail supports partial or full in rack service position
- " Requires swap screws (included in Rail Kit), based on chassis rack ear type and Rack Installation guide
- $^{\scriptscriptstyle v}$ The hole diameter of the threaded hole rack flange is required to be greater than 4mm
- $^{\scriptscriptstyle W}$ The hole diameter of the threaded hole rack flange equal or greater than 10-32UNF-2B

D&LLTechnologies





Dell Technologies Enterprise Systems Rail Sizing and Rack Compatibility Matrix Table 3. Dell Technologies rack compatibility matrix

						ll-branded APC Racks 100X717/AR3104X717)	Dell xx20	Jell xx20D/xx20S	Dell xx20W	Dell xx10	HP 10XXX	HP/Compaq 9XXX	IBM S2	APC Netshelter SX m Wide x 1070mm Deep)	'Post Rack Spacing TITAN-D/ TITAN S	' Post Rack Spacing TITAN /TITAN SS	iebert Foundation	atsworth Teraframe	ightline Vantage S2
		Product	Rail Identifier	Mounting Interface	Rail Type	De (AR3								(600m	24'.	24'		Ĉ	٨r
			Α7	ReadyRails II	Sliding	√ ²	×	×	~	×	✓	√1	×	√2	X	X	✓	×	×
		R640 (8-HDD)	A8	ReadyRails	Stab-in Static	*	*	×	<	×	<	<	<	×	✓15	√ ¹⁵	*	~	×
			A10	Generic Tool-less	Drop-in/ Stab-in	√ ²	~	×	~	×	~	✓1	×	√ 2	√14	√ 14	✓	~	~
			Α7	ReadyRails II	Sliding	√3,4	√ ²	√ 2	<	√9	<	▲	<	√3,4	x	x	*	~	×
		R640 (4-HDD/10-HDD)	A8	ReadyRails	Stab-in Static	~	~	~	1	~	*	1	*	*	√ ¹⁵	√ ¹⁵	*	~	~
	4		A10	Generic Tool-less	Drop-in/ Stab-in	√3,4	√ ²	√ ²	✓	√9	✓	✓1	✓	√ ^{3,4}	√ ¹⁴	√ ¹⁴	✓	×	~
ERS	:dge ^{TA}	R6525 (8-HDD)	A15	ReadyRails II	Sliding	√4,12	√2	√ ²	✓	✓	✓	✓1	✓	√ ^{3,4}	x	x	✓	~	~
SERV	wer	R650 (8-HDD) R660/R6615/R6625	A14	ReadyRails	Stab-in Static	×	✓	×	~	×	×	✓	×	✓	√ ¹⁵	√ ¹⁵	~	×	×
	Pc	(0-HDD/8-HDD)	A16	Generic Tool-less	Drop-in/ Stab-in	√4,12	√ ²	√ 2	✓	✓	~	√1	~	√3,4	√14	√ ¹⁴	✓	~	~
		R6525 (4-HDD/10-HDD)	A15	ReadyRails II	Sliding	√ ^{4,12}	√4,12	√4,12	*	√ ¹³	~	✓1	√ ¹³	√ ^{3,4}	x	x	~	×	√ ¹³
		R650 (4-HDD/10-HDD) R660/ R6615/R6625 (4-HDD/10-HDD/E3/LO-7)	A14	ReadyRails	Stab-in Static	×	*	×	✓	✓	✓	*	✓	✓	√ ¹⁵	√ ¹⁵	✓	~	~
		R670/R6725/R6715/R470	A16	Generic Tool-less	Drop-in/ Stab-in	√ ^{4,12}	√ ^{4,12}	√ ^{4,12}	✓	√ ¹³	~	✓1	√ ¹³	√ ^{3,4}	✓ ¹⁴	√ ¹⁴	×	~	√ ¹³
		R340 (8-HDD)	A12	ReadyRails II	Sliding	√ 3,4	√ ²	✓2	×	√9	✓	✓1	×	√3,4	X	x	✓	~	✓
		R360 (8-HDD)	A8	ReadyRails	Stab-in Static	~	✓	×	✓	×	~	~	✓	*	√ ¹⁵	√15	~	~	~

22

D&LLTechnologies



				ll-branded APC Racks 100X717/AR3104X717)	Dell xx20	əeli xx20D/xx20S	Dell xx20W	Dell xx10	HP 10XXX	HP/Compaq 9XXX	IBM S2	APC Netshelter SX n Wide x 1070mm Deep)	Post Rack Spacing ITAN-D/ TITAN S	Post Rack Spacing TITAN /TITAN SS	ebert Foundation	itsworth Teraframe	ightline Vantage S2
Product	Rail Identifier	Mounting Interface	Rail Type	Del (AR3						-		(600mr	24" T	24".	Li	Cha	Wri
R340 (4-HDD)	A12	ReadyRails II	Sliding	√3,4	√2	√ ²	×	√9	~	√1	~	√3,4	x	x	*	~	~
R360 (4-HDD)	A8	ReadyRails	Stab-in Static	*	~	~	×	*	*	*	*	×	√ ¹⁵	√ 15	*	~	~
R440 (8-HDD) R450 (8-HDD) R6415 (8-HDD)	A8	ReadyRails	Stab-in Static	~	1	1	*	~	~	~	*	~	√ ¹⁵	√ 15	~	~	~
R6515 (8-HDD) R650xs (0-HDD/8-HDD) R660xs(0-HDD/8-HDD	A11	Generic Tool-less	Drop-in/ Stab-in	✓2	*	*	*	*	*	√1	>	√ ²	√14	√14	*	*	*
R440 (4-HDD/10-HDD) R450 (4-HDD) R6415 (4-HDD/10-HDD)	A8	ReadyRails	Stab-in Static	*	*	*	*	~	*	*	*	~	√ ¹⁵	√15	*	~	~
R6515 (4-HDD/10-HDD) R650xs (4-HDD/10-HDD) R650xs (8-HDD NVME) R660xs(4-HDD/10-HDD)	A11	Generic Tool-less	Drop-in/ Stab-in	√ 3,4	√2	√2	*	√9	<	√1	*	√3,4	√14	√14	*	~	~
R540/R540xd	B6	ReadyRails II	Sliding	√ ²	×	~	~	~	~	√1	×	√2	x	X	~	~	~
R740/R740xd/ R7415/R7425/R7515	B4	ReadyRails	Stab-in Static	~	×	~	×	×	×	×	×	×	√ ¹⁵	√ 15	✓	~	~
	B13	Generic Tool-less	Drop-in/ Stab-in	✓2	~	✓	×	✓	*	✓1	✓	✓2	√ ¹⁴	√ ¹⁴	✓	~	~
R550	B21	ReadyRails II	Sliding	✓2	~	~	×	×	1	✓1	✓	√ ²	x	x	~	~	×
R750xs R760xs	B20	ReadyRails	Stab-in Static	~	~	~	 Image: A start of the start of	✓	~	×	~	✓	√ ¹⁵	√ 15	✓	~	×
	B22	Generic Tool-less	Drop-in/ Stab-in	✓2	~	×	×	~	*	✓1	×	√ ²	√14	√ ¹⁴	*	~	×

D&LLTechnologies



				ll-branded APC Racks 100X717/AR3104X717)	Dell xx20	Dell xx20D/xx20S	Dell xx20W	Dell xx10	HP 10XXX	HP/Compaq 9XXX	IBM S2	APC Netshelter SX m Wide x 1070mm Deep)	'Post Rack Spacing TITAN-D/ TITAN S	'Post Rack Spacing TITAN /TITAN SS	iebert Foundation	atsworth Teraframe	ightline Vantage S2
Product	Rail Identifier	Mounting Interface	Rail Type	De (AR3								(600m	24,	24'		Ċ	Wr
	B21	ReadyRails II	Sliding	√ 3,5,12	√2	√2	*	×	×	√1	×	√3,4	x	x	*	~	×
R760/R7615/R7625 R770/R7715/R7725/R570	B20	ReadyRails	Stab-in Static	~	1	~	×	×	×	*	×	×	√ ¹⁵	√ ¹⁵	✓	×	~
	B22	Generic Tool-less	Drop-in/ Stab-in	√ ^{3,5,12}	√ ²	√ ²	*	×	×	✓1	×	√ ^{3,4}	✓ ¹⁴	√14	*	~	×
	B6	ReadyRails II	Sliding	√ ^{3,5,12}	√ ²	√ ²	×	×	✓	✓1	✓	√ ^{3,4}	x	x	✓	~	~
R7525 R750	Β4	ReadyRails	Stab-in Static	~	~	~	×	×	✓	×	✓	×	√ ¹⁵	√ ¹⁵	✓	~	~
	B13	Generic Tool-less	Drop-in/ Stab-in	√ ^{3,5,12}	√ ²	√ ²	*	×	×	~	×	√ ^{3,4}	✓14	✓14	*	~	~
P750va	B19	Generic Tool-less	Drop-in/ Stab-in	√4,6,12	√5	√5	*	√ 10,13	√ ^{10,13}	√ ^{10,13}	√ ¹⁰	✓	√ ¹⁴	√ ¹⁴	*	√ 10,13	✓10,1 3
KY JOXA	B17	ReadyRails	Stab-in Static	√4,6	√4	√4	~	√ 10	√ ¹⁰	√ 10	√ ¹⁰	×	√ ¹⁵	√ 15	~	√ 10,13	✓10,1 3
P760va	B25	Generic Tool-less	Drop-in/ Stab-in	√4,7,12	√4,7	~	✓	√ ^{10,13}	√ 10,13	√ ^{10,13}	√ 10,13	√4,7,12	√ ¹⁴	√ 13,14	√ ¹³	√ 10,13	✓10,1 3
K7 OOXa	B33	ReadyRails	Stab-in Static	√4,7	√4,7	~	✓	√ ¹⁰	√ ¹⁰	√ ¹⁰	√ ¹⁰	√ 4,7	√ ¹⁴	√ 14	~	√ 10	√ ¹⁰
R770 Cold Aisle	B35	ReadyRails	Stab-in, Static	√4,7	√4,7	~	~	√ 10	√10	√ 10	√10	√4,7	x	x	~	√10	√10
R670 Cold Aisle	A28	ReadyRails	Stab-in, Static	√4,7	√4,7	~	✓	√ 10	√10	√ 10	√10	√4,7	x	x	~	√10	√ ¹⁰
R740xd2	-	Generic Tool-less	L-Bracket Static	~	~	✓	✓	1	1	1	✓	~	√ ¹⁴	√ 14	~	~	~

D&LLTechnologies



				ll-branded APC Racks 100X717/AR3104X717)	Dell xx20	Jell xx20D/xx20S	Dell xx20W	Dell xx10	HP 10XXX	HP/Compaq 9XXX	IBM S2	APC Netshelter SX m Wide x 1070mm Deep)	' Post Rack Spacing FITAN-D/ TITAN S	' Post Rack Spacing TITAN /TITAN SS	iebert Foundation	atsworth Teraframe	ightline Vantage S2
Product	Rail Identifier	Mounting Interface	Rail Type	De (AR3								(600m	24"]	24"		Ċ	Wr
R760xd2	B23	Generic Tool-less	Stab-in Sliding	√ 4,7,12 ,13	√ 4,7,12 13	*	√ ¹³	✓ 10,11 • 13,17	√ 10,13	√ 10,13	√ 10,13	✓4,7,12 ,13	x	x	√ ¹³	√ 10,13	✓10,1 3
R840	B15	Generic Tool-less	Drop-in/ Stab-in	√ ^{4,6,12}	√5	√5	✓	√ ¹⁰	√ ¹⁰	√ ^{10,13}	√ ¹⁰	~	x	x	✓	√ 10	√ 10,1 3
50/0	B19	Generic Tool-less	Drop-in/ Stab-in	√ 4,7,12 ,13	√7,13	×	√ ¹³	√ 10,13	√ 10,13	√ 10,13	√ 10,13	√4,7,12 ,13	√ ¹⁴	✔ 10,13,14	√ ¹³	√ ¹³	√ 10,1 3
K860	B24	ReadyRails	Stab-in Static	√4,6	√6	~	~	√ 10	√ ¹⁰	√10	√10	√4,6	√ ¹⁵	√ 15	1	~	√ 10
R940	B12	ReadyRails II	Sliding	✓3,6,12	√ ^{3,6,12}	√ 3,6,12	~	√ ¹³	√ ¹³	√ ¹³	√ ¹³	~	√ ¹⁵	√ ¹⁵	1	~	~
R960	B25	Generic Tool-less	Drop-in/ Stab-in	✓4,7,12 ,13	√7,13	~	√ ¹³	√ 10,13	√ 10,13	√ 10,13	√ 10,13	√14	✓10,13,1 4	√14	√ ¹³	√ 13	✓10,1 3
R940xa	B16	Generic Tool-less	Drop-in/ Stab-in	√ ^{4,6,12}	√5	√5	*	√ 10	√ 10	√ 10,13	√10	~	x	x	1	√ 10	✓10,1 3
EV2/EV2-	B10	ReadyRails II	Sliding	√4,6,12	✓ ⁵	✓ ⁵	~	✓10	√ ¹⁰	√ 10,13	√ ¹⁰	√4,6,12	x	x	✓	√ 10	√10,1 3
FAZ/FAZS	B11	ReadyRails II	Stab-in Static	√4,6	~	*	~	√10	√10	√ 10	√ 10	√4,6	x	x	✓	√ 10	√ 10
C4130/C4140	Α9	ReadyRails II	Stab-in Static	√7	√ 4,7,10	√ 4,7,10	√ ¹⁰	x	x	X	X	√7	X	X	√ ¹⁰	X	X
VRTX	С3	ReadyRails II	Sliding	√2	~	~	1	~	~	✓1	~	√2	√ 15	√ 15	✓	~	~
R240/R250/R260	Α4	ReadyRails	Stab-in Static	~	~	~	×	×	×	~	~	~	√ ¹⁵	√ 15	✓	~	~
M1000e	-	RapidRails	L-Bracket Static	√4,5	~	~	~	~	~	~	~	√ 4,5	X	X	1	~	~
miouce	-	VersaRails	L-Bracket Static	√4,5	✓	✓	1	✓	✓	✓	✓	√ 4,5	X	X	1	✓	×

D&LLTechnologies



					ll-branded APC Racks 100X717/AR3104X717)	Dell xx20	Jell xx20D/xx20S	Dell xx20W	Dell xx10	HP 10XXX	HP/Compaq 9XXX	IBM S2	APC Netshelter SX m Wide x 1070mm Deep)	' Post Rack Spacing FITAN-D/ TITAN S	' Post Rack Spacing TITAN /TITAN SS	iebert Foundation	itsworth Teraframe	ightline Vantage S2
	Product	Rail Identifier	Mounting Interface	Rail Type	De (AR3						-		(600m	24" 1	24"	ï	Chả	Wr
	MX7000	C5	ReadyRails II	L-Bracket Static	√ 4,6,16	√ ¹⁶	√ ¹⁶	√ ¹⁶	√ 10,16	√ ¹⁰	√ 10	√10	√ 4,6,16	√ ¹⁵	√15	✓	√10	√ ¹⁰
		A12	ReadyRails II	Sliding	√3,4	√ ²	√ ²	~	√9	~	✓1	~	√3,4	X	x	*	~	~
	XR2	A4	ReadyRails	Stab-in Static	~	~	~	~	~	×	×	✓	~	√ ¹⁵	✓ ¹⁵	*	~	~
		-	Generic	Stab-in Static ^t	-	-	-	-	-	-	-	-	-	-	-	-	-	-
	VD44 (VD42	A20	Generic Tool-less	Stab-in Sliding	~	✓	~	~	~	×	×	✓	~	√ ¹⁴	√14	1	~	~
	XR11/XR12	A21 ¹⁸	Generic Tool-less	Stab-in Sliding	-	-	-	-	-	-	-	-	-	-	-	-	-	-
	C1100	-	Generic Tool-less	Sliding	~	✓	~	✓	×	×	×	✓	×	x	x	✓	~	~
	C2100	-	Generic	Sliding	~	✓	~	✓	~	✓	×	✓	~	X	x	~	~	1
	C410x	-	VersaRails	Stab-in Static	√ ⁸	√ ⁸	√ ⁸	√ ⁸	√ ⁸	1	*	√ ⁸	√ 8	X	x	x	X	~
lge C	C5xxx	-	Generic Tool-less	L-Bracket Static	×	*	~	*	*	×	×	*	×	X	x	✓	~	~
PowerE	C64xx C65xx C66xx	-	Generic Tool-less	L-Bracket Static	*	1	~	1	√16	*	*	1	*	√14	√14	*	~	~
		-	Generic Tool-less	L-Bracket Static	√4	✓	~	✓	~	✓	×	✓	√4	X	X	✓	~	~
	C8000	-	Generic Tool-less	Sliding	√4,6	√ ^{4, 11}	√ 4, 11	√11	*	1	*	1	√4,6	x	x	~	~	~

D&LLTechnologies



					ll-branded APC Racks 100X717/AR3104X717)	Dell xx20	Jell xx20D/xx20S	Dell xx20W	Dell xx10	HP 10XXX	HP/Compaq 9XXX	IBM S2	APC Netshelter SX m Wide x 1070mm Deep)	' Post Rack Spacing FITAN-D/ TITAN S	' Post Rack Spacing TITAN /TITAN SS	iebert Foundation	atsworth Teraframe	ightline Vantage S2
	Product	Rail Identifier	Mounting Interface	Rail Type	De (AR3								(600m	24'	24'		Chi	Wr
		B21	ReadyRails II	Sliding	√2	~	~	~	~	~	√1	1	√2	x	x	~	~	~
	HS5620	B20	ReadyRails	Stab-in Static	~	~	~	~	~	*	~	1	~	√ ¹⁵	√15	*	~	~
SP		B22	Generic Tool-less	Drop-in/ Stab-in	√2	~	*	~	×	×	✓1	1	√ 2	√ ¹⁴	√14	*	×	~
Edge C	HS5610 Cold Aisle	A22	ReadyRails	Stab-in, Static	~	~	~	~	×	*	*	✓	~	√ ¹⁵	√ 15	~	×	~
ower		A8	ReadyRails	Stab-in Static	~	1	~	~	×	✓	×	1	×	√15	√ ¹⁵	✓	✓	~
₽.	133010(0-0007)	A11	Generic Tool-less	Drop-in/ Stab-in	√ ²	*	~	×	×	×	✓1	*	√ ²	√14	√ ¹⁴	✓	×	~
		A8	ReadyRails	Stab-in Static	~	✓	~	~	×	✓	×	*	×	√ ¹⁵	√ ¹⁵	✓	×	~
	HS5610(4-HDD/10-HDD)	A11	Generic Tool-less	Drop-in/ Stab-in	√3,4	√ ²	√2	~	√9	*	√1	1	√ ^{3,4}	√ ¹⁴	√14	*	~	~
		A23	Generic	Stab-in	~	~	~	~	×	*	*	✓	×	*	~	<	*	~
	XR4000R	A24 ¹⁸	Tool-less	Sliding	X	X	X	X	X	X	X	X	X	X	X	X	X	X
×		A25 ¹⁸			x	x	x	x	X	x	x	X	X	x	x	x	x	x
rEdge	YPSOOP	B31 ¹⁸	Generic Tool-less	L-Bracket Static	X	X	X	X	X	X	X	X	X	X	X	X	X	X
Power		B36 ¹⁸	Generic Tool-less	L-Bracket Static	x	X	x	X	X	X	X	X	X	X	x	x	X	X
	VR5610	A26	Generic	Stab-in	~	~	~	~	1	×	1	~	1	~	~	~	×	~
	υΓόςηλ	A27 ¹⁸	Tool-less	Sliding	x	X	x	X	X	X	x	X	X	x	x	х	x	X

D&LLTechnologies



Dell Technologies Enterprise Systems Rail Sizing and Rack Compatibility Matrix

					ll-branded APC Racks 100X717/AR3104X717)	Dell xx20	0ell xx20D/xx20S	Dell xx20W	Dell xx10	HP 10XXX	HP/Compaq 9XXX	IBM S2	APC Netshelter SX m Wide x 1070mm Deep)	' Post Rack Spacing FITAN-D/ TITAN S	' Post Rack Spacing TITAN /TITAN SS	iebert Foundation	itsworth Teraframe	ightline Vantage S2
	Product	Rail Identifier	Mounting Interface	Rail Type	De (AR3						-		(600m	24" 1	24"	5	Cha	Wr
	XR7620	B29	Generic	Stab-in	~	~	~	×	~	×	×	~	~	×	~	~	×	~
	ART 020	B30 ¹⁸	Tool-less	Sliding	X	X	X	X	X	X	X	X	X	X	X	X	X	x
	XE2420	-	Generic Tool-less	Stab-in Static Standard	√17	√17	√17	√17	√17	√17	√17	√17	√17	x	x	√17	√17	√17
	XE8640	B27	Generic Tool-less	Stab-in Sliding	√6,12	√7	~	~	√13	√13	√13	√ ¹³	√6,13	√14	√13,14	*	√13	√13
XE	XE9680	B28	Generic Tool-less	Stab-in Sliding	x	x	√4,12	x	x	x	x	x	x	√ 4,12, 14	x	x	x	x
PowerEdge	XE9640	B25	Generic Tool-less	Drop-in/ Stab-in	√4,7,12 ,13	√4,7	~	~	√10,13	√ 10,13	√ 10,13	√10,13	√ 4,7,12 ,13	√14	√13,14	*	√10,13	√ 10,1 3
	XE9680L/XE9685L	B38	Generic Tool-less	Stab-in Sliding	x	x	√4,7	x	x	x	x	x	x	√4,14	x	x	x	x
	XE7740/XE7745	B37	Generic Tool-less	Stab-in Sliding	x	x	√4,6	x	x	x	x	x	x	√4,14	x	x	x	x

28

D&LLTechnologies



Dell Technologies Enterprise Systems Rail Sizing and Rack Compatibility Matrix

						ll-branded APC Racks 100X717/AR3104X717)	Dell xx20)ell xx20D/xx20S	Dell xx20W	Dell xx10	HP 10XXX	HP/Compaq 9XXX	IBM S2	APC Netshelter SX n Wide x 1070mm Deep)	Post Rack Spacing TITAN-D/ TITAN S	Post Rack Spacing TITAN /TITAN SS	ebert Foundation	itsworth Teraframe	ightline Vantage S2
		Product	Rail Identifier	Mounting Interface	Rail Type	Del (AR3						-		(600m)	24" 1	24".	Li.	Cha	Wri
	٨	1081AD/2161AD/1082DS /2162DS 4322DS	Α5	ReadyRails	Stab-in Static	~	~	~	~	~	*	*	~	*	-	-	~	*	~
	KVN	180AS/2160AS 2161DS/2161DS- 2/4161DS 2321DS	-	Generic	Stab-in Static	*	*	*	*	*	~	*	*	*	-	-	*	*	*
ITCHES		PC8132/PC8132F	Α5	ReadyRails	Stab-in Static	*	*	*	*	*	*	<	<	*	-	-	*	*	✓
SWI	ing	PC8164/PC8164F																	
	Networki	S4820T/S6000	А5	ReadyRails	Stab-in Static	~	~	~	~	*	*	*	*	*	-	-	~	*	*
		\$5000	-	Generic	Stab-in Static	~	*	~	~	~	*	*	~	~	-	-	~	~	~
		R7910	B6	ReadyRails II	Sliding	✓2	*	~	×	×	~	√1	×	√ ²	-	-	×	✓	~
			B4	ReadyRails	Stab-in Static	*	~	1	~	×	~	<	<	×	-	-	×	✓	×
	SNOI	Precision 3930 Rack	A4	ReadyRails	Stab-in Static	×	×	~	×	×	×	×	×	×	-	-	✓	✓	~
	STAT		B6	ReadyRails II	Sliding	√ ²	×	~	×	×	✓	✓1	×	√ ²	x	x	✓	~	×
	WORK		B4	ReadyRails	Stab-in Static	~	✓	✓	 	×	×	×	×	×	√15	√ 15	✓	✓	✓
		Precision 7920 Rack	B13	Generic Tool-less	Drop-in/ Stab-in	√ ²	~	*	~	~	*	√1	*	√ ²	√14	√14	*	*	~

29

D&LLTechnologies



					l-branded APC Racks 100X717/AR3104X717)	Dell xx20	ell xx20D/xx20S	Dell xx20W	Dell xx10	HP 10XXX	IP/Compaq 9XXX	IBM S2	APC Netshelter SX n Wide x 1070mm Deep)	Post Rack Spacing ITAN-D/ TITAN S	Post Rack Spacing TITAN /TITAN SS	ebert Foundation	tsworth Teraframe	ghtline Vantage S2
	Product	Rail Identifier	Mounting Interface	Rail Type	Del (AR3 1		Δ				-		лтоод) (24" T	24"]	Li	Cha	Wri
		B21	ReadyRails II	Sliding	√ 3,5,12	√2	√ ²	×	*	*	✓1	×	√3,4	X	x	*	×	~
	Precision 7960 Rack	B20	ReadyRails	Stab-in Static	1	~	~	~	~	~	×	×	×	√ ¹⁵	√ ¹⁵	*	×	~
		B22	Generic Tool-less	Drop-in/ Stab-in	√ 3,5,12	√ ²	√ ²	~	✓	~	✓1	✓	√ 3,4	√14	√ 14	*	×	~
	T7600/T7610	C2	ReadyRails II	Sliding	✓2	✓11	✓11	✓11	*	*	✓1	×	✓2	-	-	*	×	~
	R5500/R7610	B2	ReadyRails	Sliding	√3	√2	✓2	~	~	~	✓1	×	✓3	-	-	*	×	~
	FPM185 (without KVM)		ReadyRails II	Sliding	~	*	*	~	~	~	*	~	~	-	-	x	×	~
KMM	FPM185 (with KVM)	-	ReadyRails II	Sliding	~	~	~	~	~	~	~	~	~	-	-	x	~	~
	1750	-	RapidRails	Sliding	~	~	~	~	×	×	×	×	×	-	-	✓	*	~
	1766	-	VersaRails	Sliding	~	~	~	~	*	*	~	~	*	-	-	*	*	✓
NPS	Dell Rack Mount UPS Family	B5	ReadyRails	Stab-in Static	~	~	~	~	*	*	*	*	*	-	-	~	*	×
OTHER	1U Fixed Equipment Shelf	Α4	ReadyRails	Stab-in Static	~	*	~	~	*	*	*	*	*	-	-	~	*	*
STORAGE	PROXPE1 PROXPE2	B14	-	L-Bracket Static	√ 6,13,1 6	*	*	✓13	√13	√13	√13	√ 13	~	~	*	√13	√13	√13

D&LLTechnologies



					ll-branded APC Racks 100X717/AR3104X717)	Dell xx20	Jell xx20D/xx20S	Dell xx20W	Dell xx10	HP 10XXX	HP/Compaq 9XXX	IBM S2	APC Netshelter SX m Wide x 1070mm Deep)	'Post Rack Spacing TITAN-D/ TITAN S	'Post Rack Spacing TITAN /TITAN SS	iebert Foundation	atsworth Teraframe	ightline Vantage S2
	Product	Rail Identifier	Mounting Interface	Rail Type	De (AR3								(600m	24'.	24'		Chi	Wr
	NX3300/NX400	Α7	ReadyRails II	Sliding	√ ²	~	×	×	×	*	✓	*	X	X	X	*	✓	~
	NX3300/NX400	A8	ReadyRails	Stab-in Static	~	~	×	×	×	×	*	*	√ ¹⁵	√ ¹⁵	√ 15	*	✓	✓
	NY3200	B6	ReadyRails II	Sliding	√ ²	~	~	×	×	×	✓1	✓	x	X	X	✓	×	✓
	NAS200	B4	ReadyRails	Stab-in Static	~	~	~	×	×	×	×	✓	√ ¹⁵	√ ¹⁵	√ ¹⁵	✓	×	✓
		A3	ReadyRails	Sliding	✓2	~	~	*	*	*	✓1	~	x	x	x	*	*	~
	NX3500 Controller	A4	ReadyRails	Stab-in Static	~	~	✓	✓	✓	×	×	✓	√ ¹⁵	√ ¹⁵	√ ¹⁵	1	×	~
אמחור	NX3500 UPS	A4	ReadyRails	Stab-in Static	~	~	~	~	~	~	*	*	~	√ ¹⁵	√ 15	~	~	~
	DX6000G	A4	ReadyRails	Stab-in Static	~	~	~	~	~	~	*	*	~	√ 15	√15	~	✓	~
	DABOOOG	A6	ReadyRails	Stab-in Static	~	~	~	×	*	*	*	*	*	X	X	*	*	×
	NX300/DX6004S	A3	ReadyRails	Sliding	✓2	~	~	×	×	×	✓1	~	✓2	x	x	~	×	~
		A4	ReadyRails	Stab-in Static	~	~	~	✓	✓	✓	✓	~	✓	√ ¹⁵	√ 15	~	×	1
	NX3000/DX6000	B1	ReadyRails	Sliding	✓2	~	~	✓	✓	✓	✓1	~	√ ²	X	x	✓	✓	~
		A2	ReadyRails	Stab-in Static	~	~	~	✓	~	✓	~	~	✓	√ ¹⁵	✓15	✓	✓	~
	NX3100/DL2200/	В3	ReadyRails	Sliding	✓2	~	~	~	~	~	✓1	~	✓2	x	x	*	*	~

D&LLTechnologies



				ll-branded APC Racks 100X717/AR3104X717)	Dell xx20	Jell xx20D/xx20S	Dell xx20W	Dell xx10	HP 10XXX	HP/Compaq 9XXX	IBM S2	APC Netshelter SX m Wide x 1070mm Deep)	' Post Rack Spacing FITAN-D/ TITAN S	 Post Rack Spacing TITAN /TITAN SS 	iebert Foundation	itsworth Teraframe	ightline Vantage S2
Product	Rail Identifier	Mounting Interface	Rail Type	De (AR3						-		(600m	24" 1	24"	5	Chĕ	Wr
DX6012S/DR4000	B4	ReadyRails	Stab-in Static	~	~	~	×	×	×	×	~	×	√ ¹⁵	✓15	×	~	 Image: A start of the start of
MD3060e/MD3660	-	VersaRails	L-Bracket Static	√4,6	√4	✓4	×	✓10	✓10	x	√ 10	√4,6	x	x	✓	√10	X
MD12xx/14xx/32xx/36xx	В9	ReadyRails II	L-Bracket Static	~	1	~	*	*	*	×	~	*	✓15	✓15	✓	~	~
NV26vv	-	RapidRails	L-Bracket Static	~	~	~	~	×	*	x	~	×	x	x	x	~	~
NAJOXX	-	VersaRails	L-Bracket Static	~	~	×	×	×	×	×	*	×	X	x	×	~	×
MD1120	-	RapidRails	L-Bracket Static	~	*	*	×	×	*	X	*	×	x	x	X	~	✓
MD1120	-	VersaRails	L-Bracket Static	×	*	×	×	×	×	×	*	×	x	x	×	~	✓
MD1000/MD3000	-	RapidRails	L-Bracket Static	×	✓	~	×	×	✓	×	×	×	x	x	✓	~	~
MD 1000/MD3000	-	VersaRails	L-Bracket Static	✓	*	*	×	×	*	*	*	×	x	x	✓	~	~
	Β7	ReadyRails	Stab-in Static	~	~	*	~	*	~	*	*	*	√ ¹⁵	✓15	~	~	~
PV114T/PV114X	-	RapidRails	Sliding	√2	~	~	×	×	*	✓1	~	✓2	x	x	✓	~	~
	-	VersaRails	Sliding	✓2	~	~	~	~	×	✓1	~	✓2	X	X	~	~	~
	-	RapidRails	L-Bracket Static	~	~	~	~	~	~	~	~	~	X	X	~	~	~
PV124T	-	VersaRails	L-Bracket Static	*	~	~	*	*	*	*	~	~	x	x	*	~	~

D&LLTechnologies



						ll-branded APC Racks 100X717/AR3104X717)	Dell xx20	Jell xx20D/xx20S	Dell xx20W	Dell xx10	HP 10XXX	HP/Compaq 9XXX	IBM S2	APC Netshelter SX m Wide x 1070mm Deep)	' Post Rack Spacing FITAN-D/ TITAN S	' Post Rack Spacing TITAN /TITAN SS	iebert Foundation	atsworth Teraframe	ightline Vantage S2
		Product	Rail Identifier	Mounting Interface	Rail Type	De (AR3								(600m)	24" 1	24"	Ē	ĊŸ	Wr
		ES7E00 Controllor	A1	ReadyRails	Sliding	√3	√ ²	√ ²	✓	~	*	√1	✓	√3	x	x	~	~	~
		F37500 Controller	A2	ReadyRails	Stab-in Static	~	*	*	✓	*	✓	*	*	~	√15	√15	*	~	~
	alLogic	FS7500 UPS	A4	ReadyRails	Stab-in Static	~	*	*	*	*	*	~	*	*	√ 15	✓15	*	~	~
		FS76xx/PS41xx/PS61xx	В9	ReadyRails II	L-Bracket Static	~	~	×	~	~	~	~	×	×	√ ¹⁵	√ 15	~	~	~
			-	RapidRails	L-Bracket Static	~	~	×	✓	✓	✓	x	*	×	x	x	x	~	~
	Equ		-	VersaRails	L-Bracket Static	~	~	~	~	~	~	✓	✓	~	x	x	*	~	~
		PS6500/6510	-	ReadyRails	Sliding	√7	√2	√2	*	*	*	x	*	√7	√ 15	√15	*	~	~
		PS4000/6000/6010	-	Generic	L-Bracket Static	~	~	~	~	*	*	✓	~	~	x	x	*	*	~
	Dell Compellent	SC20xx/SC40xx	-	Generic	L-Bracket Static	~	~	*	*	*	*	*	*	*	√15	√15	*	*	*

D&LLTechnologies



		_		ll-branded APC Racks 100X717/AR3104X717)	Dell xx20	Dell xx20D/xx20S	Dell xx20W	Dell xx10	HP 10XXX	HP/Compaq 9XXX	IBM S2	APC Netshelter SX m Wide x 1070mm Deep)	'Post Rack Spacing TITAN-D/ TITAN S	'Post Rack Spacing TITAN /TITAN SS	iebert Foundation	atsworth Teraframe	ightline Vantage S2
Product	Rail Identifier	Mounting Interface	Rail Type	De (AR3								(600m	24'	24'		Ċ	Wr
	B6	ReadyRails II	Sliding	√2	*	✓	~	*	*	√1	*	√ ²	√ ¹⁵	√ 15	✓	✓	~
SC8000	B4	ReadyRails	Stab-in Static	*	*	1	*	*	*	*	*	*	√15	√15	*	✓	*
	В9	ReadyRails II	L-Bracket Static	~	×	~	~	~	~	1	1	×	x	x	~	~	~
SC2xx/FS86xx	-	RapidRails	L-Bracket Static	×	×	×	×	*	*	X	×	×	x	x	x	~	~
	-	VersaRails	L-Bracket Static	*	~	~	*	*	*	~	<	*	x	x	*	~	~
SCV30xx/SC50xx/SC	С7020 В9	ReadyRails II	L-Bracket Static	~	✓	~	~	~	×	*	*	×	√ ¹⁵	√ 15	×	~	~
Series 40	-	Generic	Sliding	~	~	~	~	~	~	*	*	*	~	~	✓	~	~
Fibre Channel	-	Generic	Stab-in Static	×	~	×	×	*	*	~	~	×	~	×	✓	~	~
SAS (new rails) -	Generic	Stab-in Static	*	1	~	*	*	*	×	~	×	X	x	~	~	~
SAS (old rails)	-	Generic	Stab-in Static	~	✓	~	~	~	~	×	~	~	X	X	✓	✓	~
NAS Gen3	-	Generic	Sliding	√ 6	×	~	×	×	×	×	×	√6	X	x	~	✓	~

D&LLTechnologies



Notes:

- ¹ A rear door extension kit is required to accommodate the CMA.
- ² CMA may impede access to forward bank of rear-mount PDUs.
- ³ CMA and outer CMA brackets must be removed in order to access the forward bank of rear-mount PDUs.
- ⁴ Rear-mount PDUs may impede extraction of some rear system modules.
- ⁵ The strain relief bar interferes with the forward bank of rear-mount PDUs.
- ⁶ Rails/system block the forward bank of rear-mount PDUs.
- ⁷ Rails/system block both the forward and rearward banks of rear-mount PDUs.
- ⁸ The rear mounting flanges of the rack must be moved rearward.
- $^{\rm 9}~$ The CMA tray interferes with rear door lock rod in top U and bottom U.
- ¹⁰ Space for external cable routing is limited.
- ¹¹ May need to adjust the rack's mounting posts back to allow the front door to close.
- ¹² CMA/SRB fully blocks front bank of rear-mount PDUs, and partially blocks the rearward PDU banks. Recommend rotating PDUs 90°.
- ¹³ CMA/SRB must be removed to enable rear door to close for some or all racks in this column
- ¹⁴ The rails align with bezels on Storage systems (unthreaded round-hole rack).
- ¹⁵ The rails require tooled installation for bezel alignment with Storage systems (unthreaded round-hole rack).
- ¹⁶ Strain relief bar might block a portion of the rearward bank of the rear-mount PDUs.
- ¹⁷ Normal Inner rail member allows for tool-less bezel installation and does not enable front rack door to close.
- ¹⁸ Rail is only intended for use within ruggedized transit case (Pelican custom rack 25-036329-01)
- ¹⁹ Rail is only intended for telco short depth rack.

D&LLTechnologies



Integrated Dell Remote Access Controller 9 User's Guide

December 2020 Rev. A02





Notes, cautions, and warnings

(i) NOTE: A NOTE indicates important information that helps you make better use of your product.

CAUTION: A CAUTION indicates either potential damage to hardware or loss of data and tells you how to avoid the problem.

WARNING: A WARNING indicates a potential for property damage, personal injury, or death.

© 2020 Dell Inc. or its subsidiaries. All rights reserved. Dell, EMC, and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be trademarks of their respective owners.

Chapter 1: Overview of iDRAC	16
Benefits of using iDRAC	
Key features	17
New features added	19
Firmware version 4.40.00.00	
Firmware version 4.30.30.30.	20
Firmware version 4.20.20.20.	21
Firmware version 4.10.10.10.	21
Firmware version 4.00.00.00	22
How to use this guide	23
Supported web browsers	
Supported OS and Hypervisors	23
iDRAC licenses	23
Types of licenses	
Methods for acquiring licenses	24
Acquiring license key from Dell Digital Locker	
License operations	
Licensed features in iDRAC9	
Interfaces and protocols to access iDRAC	
iDRAC port information	
Other documents you may need	
Contacting Dell	
Accessing documents from Dell support site	
Accessing Redfish API Guide	
Chapter 2: Logging in to iDRAC	
Force Change of Password (FCP)	
Logging into iDRAC using OpenID Connect	
Logging in to iDRAC as local user, Active Directory user, or LDAP user	
Logging in to iDRAC as a local user using a smart card	
Logging in to iDRAC as an Active Directory user using a smart card	40
Logging in to iDRAC using Single Sign-On	40
Logging in to iDRAC SSO using iDRAC web interface	40
Logging in to iDRAC SSO using CMC web interface	40
Accessing iDRAC using remote RACADM	41
Validating CA certificate to use remote RACADM on Linux	41
Accessing iDRAC using local RACADM	41
Accessing iDRAC using firmware RACADM	41
Simple 2-Factor Authentication (Simple 2FA)	41
RSA SecurID 2FA	42
Viewing system health	
Logging in to iDRAC using public key authentication	43
Multiple iDRAC sessions	
Secure default password	44



Resetting default iDRAC password locally	
Resetting default iDRAC password remotely	
Changing the default login password	
Changing the default login password using web interface	
Changing the default login password using RACADM	
Changing the default login password using iDRAC settings utility	
Enabling or disabling default password warning message	
Password Strength Policy	
IP Blocking	
Enabling or disabling OS to iDRAC Pass-through using web interface	
Enabling or disabling alerts using RACADM	
Chapter 3: Setting up managed system	50
	50
	51 آ
Setting up iDRAC IP using the CMC web interface	51
Auto discovery	
Configuring conversional conversion components using Auto Config	
Using hash passwords for improved security	
Modifying local administrator account sottings	01 63
Setting up managed system location	
Setting up managed system location using web interface	
Setting up managed system location using RACADM	
Setting up managed system location using iNRAC settings utility	63
Ontimizing system performance and power consumption	64
Modifying thermal settings using iDRAC web interface	64
Modifying thermal settings using RACADM	66
Modifying thermal settings using iDRAC settings utility	69
Modifying PCIe airflow settings using iDRAC web interface	70
Setting up management station	
Accessing iDRAC remotely	70
Configuring supported web browsers.	
Configuring Internet Explorer	
Configuring Mozilla Firefox	
Configuring web browsers to use virtual console	
Viewing localized versions of web interface	
Updating device firmware	
Updating firmware using iDRAC web interface	
Scheduling automatic firmware updates	
Updating device firmware using RACADM	
Updating firmware using CMC web interface	
Updating firmware using DUP	
Updating firmware using remote RACADM	
Updating firmware using Lifecycle Controller Remote Services	
Updating CMC firmware from iDRAC	
Viewing and managing staged updates	
Viewing and managing staged updates using iDRAC web interface	
Viewing and managing staged updates using RACADM	
Rolling back device firmware	
Rollback firmware using iDRAC web interface	

	FIS
	A MO
Rollback firmware using CMC web interface	
Rollback firmware using RACADM	85
Rollback firmware using Lifecycle Controller	
Rollback firmware using Lifecycle Controller-Remote Services	
	86
Monitoring iDRAC using other Systems Management tools	86
Support Server Configuration Profile — Import and Export	
Importing server configuration profile using iDRAC web interface	
Exporting server configuration profile using IDRAC web interface	
Secure Boot Configuration from BIOS Settings or F2	
BIOS recovery	
napter 4: Configuring iDRAC	
Viewing iDRAC information	
Viewing iDRAC information using web interface	
Viewing iDRAC information using RACADM	
Modifying network settings.	
Modifying network settings using web interface	92
Modifying network settings using local RACADM	92
Configuring IP filtering	93
Cinhar suite selection	90 QZ
Configuring cinher suite selection using iDRAC web interface	α <i>Λ</i>
Configuring cipher suite selection using PACADM	
EIPS mode	
Enabling EIDS Mode	
Ellabiling FIFS Mode	
Configuring services	
Configuring services using web interface	
Configuring services using RACADIVI	
SERVI Functionalities	
Using VINC client to manage remote server	
Configuring VNC server using RACADM	
Setting up VNC viewer with SSL encryption	
Setting up VNC viewer without SSL encryption	
Configuring front panel display	
Configuring LCD setting	
Configuring system ID LED setting	
Configuring time zone and NTP	102
Configuring time zone and NTP using iDRAC web interface	102
Configuring time zone and NTP using RACADM	
Setting first boot device	
Setting first boot device using web interface	102
Setting first boot device using RACADM	103
Setting first boot device using virtual console	
Enabling last crash screen	
Enabling or disabling OS to iDRAC Pass-through	103
Supported cards for OS to iDRAC Pass-through	
Supported operating systems for USB NIC	

OTOCO

Enabling or disabling OS to IDRAC Pass-through using RACADM		
Enabling or disabling OS to IDRAC Pass-through using RACADM	aling OS to iDRAC Pass-through using web interface	105
Enabling or disabling OS to IDRAC Pass-through using IDRAC settings utility	aling OS to iDRAC Pass-through using RACADM	106
Obtaining certificates 100 SSL server certificates signing request. 100 Automatic Certificate Enrollment. 100 Uploading server certificate 100 Viewing server certificates. 101 Downloading custom signing certificate. 101 Deleting custom SSL certificate signing certificate. 111 Deleting custom SSL certificate signing certificate. 111 Configuring multiple IDRACs using RACADM. 111 hapter 5: Delegated Authorization using OAuth 2.0. 112 hapter 5: Delegated Authorization using OAuth 2.0. 113 hapter 5: Delegated Authorization using OAuth 2.0. 114 Viewing managed system health and properties. 114 Viewing system information. 114 Viewing sensor information. 114 Viewing performance index of CPU, memory, and input output modules using web interface. 115 Monitoring performance index of CPU, memory, and input output modules using web. 116 GPU (Accelerators) Management. 122 Viewing historical temperature data 122 Viewing historical temperature data 122 Viewing network interfaces available on host OS. 122	bling OS to iDRAC Pass-through using iDRAC settings utility	106
SSL server certificates 100 Generating a new certificate signing request. 100 Automatic Certificate Enrollment. 100 Uploading server certificate. 100 Uploading server certificate. 100 Downloading custom SSL certificate signing certificate. 110 Downloading custom SSL certificate signing certificate. 111 Configuring multiple iDRACs using RACADM. 111 Disabiling access to modify iDRAC configuration settings on host system. 112 repter 5: Delegated Authorization using OAuth 2.0. 113 repter 5: Viewing iDRAC and managed system information. 114 Viewing system health and properties. 114 Viewing system inventory. 115 Viewing system inventory. 116 Viewing performance index of CPU, memory, and input output modules. 117 Monitoring performance index of CPU, memory, and input output modules using RACADM. 116 Belester toro: 118 Monitoring performance index of CPU, memory, and input output modules using RACADM. 116 Belestering the system for Fresh Air compliance. 122 Viewing historical temperature data 122 Viewing network interfaces avai		107
Generating a new certificate signing request	ficates	107
Automatic Certificate Enrollment. 106 Uploading server certificate. 107 Uploading custom SSL certificate signing certificate. 110 Downloading custom SSL certificate signing certificate. 111 Deleting custom SSL certificate signing certificate. 111 Disabiling access to modify iDRAC and managed system information. 112 apter 5: Delegated Authorization using OAuth 2.0. 113 apter 5: Viewing iDRAC and managed system information. 114 Configuring Asset Tracking. 114 Configuring Asset Tracking. 114 Monitoring performance index of CPU, memory, and input output modules. 111 Monitoring performance index of CPU, memory, and input output modules using RACADM. 116 GPU (Accelerators) Management. 116 Hoekking historical temperature data 117 Monitoring performance index for of CPU, memory, and input output modules using RACADM. 117 GPU (Accelerators) Management. 116 Uviewing historical temperature data 1	w cartificate signing request	108
Uploading server certificate 105 Uploading custom signing certificate 111 Uploading custom SSL certificate signing certificate 111 Devindading custom SSL certificate signing certificate 111 Deleting custom SSL certificate signing certificate 112 Papter 5: Delegated Authorization using OAuth 2.0 113 Papter 5: Viewing iDRAC and managed system information 114 Viewing managed system health and properties 114 Viewing system inventory 118 Viewing system information 114 Monitoring performance index of CPU, memory, and input output modules using web interface. 116 Monitoring performance index for of CPU, memory, and input output modules using RACADM 116 GPU (Accelerators) Management	ficate Enrollment	100
Opboding server certificate 110 Uploading custom signing certificate 111 Downloading custom SSL certificate signing certificate 111 Deleting custom SSL certificate signing certificate 111 Configuring multiple iDRACs using RACADM. 111 Disabiling access to modify iDRAC configuration settings on host system 111 apter 5: Delegated Authorization using OAuth 2.0 112 apter 5: Configuring asset Tracking. 114 Viewing managed system hand properties. 114 Viewing system inventory. 114 Viewing system inventory. 114 Viewing system inventory. 116 Monitoring performance index of CPU, memory, and input output modules. 117 Monitoring performance index of CPU, memory, and input output modules using RACADM. 116 Checking the system for Fresh Air compliance. 112 Viewing historical temperature data using IDRAC web interface. 122 Viewing network interfaces available on host OS using RACADM. 122 Viewing network interfaces available on host OS using RACADM. 122 Viewing network interfaces available on host OS using RACADM. 122 Viewing network interfaces available on host OS using RACADM.	r certificate	100
Versing as version and the set of t		110
Opbidding custom SQL certificate signing certificate 110 Downloading custom SSL certificate signing certificate 111 Configuring multiple IDRACs using RACADM	m signing contificate	110
Deleting custom SJL certificate signing certificate. 11 Configuring multiple iDRACs using RACADM. 11 Disabling access to modify iDRAC configuration settings on host system. 111 apter 5: Delegated Authorization using OAuth 2.0. 111 apter 5: Viewing iDRAC and managed system information. 112 configuring managed system health and properties. 114 Configuring Asset Tracking. 114 Viewing system inventory. 116 Weiving sensor information. 114 Monitoring performance index of CPU, memory, and input output modules using web interface. 117 Monitoring performance index of CPU, memory, and input output modules using RACADM. 118 GPU (Accelerators) Management. 116 GPU (Accelerators) Management. 112 Viewing historical temperature data using iDRAC web interface. 112 Viewing historical temperature data using iDRAC web interface. 112 Viewing network interfaces available on host OS using web interface. 112 Viewing network interfaces available on host OS using web interface. 112 Viewing network interfaces available on host OS using web interface. 112 Viewing network interfaces available on host OS using RACADM. 122	nti signing certificate signing cortificate	110
Deleting Outcom SSL Certificate signing Certificates 11 Configuring multiple IDRACs using RACADM 11 Disabling access to modify iDRAC configuration settings on host system 112 apter 5: Delegated Authorization using OAuth 2.0 111 repter 5: Viewing iDRAC and managed system information 114 Viewing managed system health and properties 114 Configuring Asset Tracking 114 Viewing system inventory 116 Viewing system inventory 116 Monitoring performance index of CPU, memory, and input output modules using web interface 117 Monitoring performance index of CPU, memory, and input output modules using RACADM 116 Idle Server Detection 116 GPU (Accelerators) Management 117 Configuring historical temperature data 118 GPU (Accelerators) Management 112 Viewing historical temperature data 112 Viewing instorical temperature data 112 Viewing network interfaces available on host OS 112 Viewing network interfaces available on host OS 112 Viewing retwork interfaces available on host OS 112 Viewing network interfaces available on host OS	Storn SSL certificate signing certificate	HU
Conjuguring multiple IDRACs using PACADM. 11 Disabling access to modify IDRAC configuration settings on host system. 112 Hapter 5: Delegated Authorization using OAuth 2.0. 113 Hapter 5: Viewing iDRAC and managed system information. 114 Viewing managed system health and properties. 114 Configuring Asset Tracking. 114 Viewing system inventory. 114 Viewing system inventory. 115 Monitoring performance index of CPU, memory, and input output modules. 117 Monitoring performance index of CPU, memory, and input output modules using web interface. 116 Monitoring performance index for of CPU, memory, and input output modules using RACADM. 116 Checking the system for Fresh Air compliance. 116 Checking the system for Fresh Air compliance. 112 Viewing historical temperature data using iDRAC web interface. 112 Viewing network interfaces available on host OS 112 Viewing network interfaces available on host OS using RACADM. 122 Viewing network interfaces available on host OS using RACADM. 122 Viewing network interfaces available on host OS using RACADM. 122 Viewing network interfaces available on host OS using RACADM. <td< td=""><td>SSL certificate signing certificate</td><td></td></td<>	SSL certificate signing certificate	
Disabiling access to modify IDRAC configuration settings on nost system. 11 hapter 5: Delegated Authorization using OAuth 2.0. 11 repter 5: Viewing iDRAC and managed system information. 11 Viewing system health and properties. 11 Configuring Asset Tracking. 11 Viewing system inventory. 11 Wointoring performance index of CPU, memory, and input output modules. 11 Monitoring performance index of CPU, memory, and input output modules using RACADM. 11 Monitoring performance index for of CPU, memory, and input output modules using RACADM. 11 GPU (Accelerators) Management. 11 Checking the system for Fresh Air compliance. 12 Viewing historical temperature data 12 Viewing network interfaces available on host OS 122 Viewing network interfaces available on host OS using RACADM. 122 Viewing represente card fabric connections. 123 Viewing Stores available on host OS using RACADM. 123 Viewing network interfaces available on host OS using RACADM. 123 Viewing network interfaces available on host OS using RACADM. 123 Viewing network interfaces available on host OS using RACADM. 124 Vi		
apter 5: Delegated Authorization using OAuth 2.0	modify IDRAC configuration settings on host system	112
https://www.communication	d Authorization using OAuth 2.0	113
Viewing managed system health and properties. 114 Configuring Asset Tracking. 114 Viewing system inventory. 118 Wointoring performance index of CPU, memory, and input output modules. 117 Monitoring performance index of CPU, memory, and input output modules using web interface. 118 Monitoring performance index of CPU, memory, and input output modules using RACADM. 118 Idle Server Detection. 116 GPU (Accelerators) Management. 116 Checking the system for Fresh Air compliance. 122 Viewing historical temperature data 121 Viewing historical temperature data using iDRAC web interface. 122 Viewing network interfaces available on host OS 122 Viewing network interfaces available on host OS 122 Viewing network interfaces available on host OS using RACADM. 122 Viewing network interfaces available on host OS using RACADM. 123 Viewing network interfaces available on host OS using RACADM. 124 Viewing network interfaces available on host OS using RACADM. 125 Viewing network interfaces available on host OS using RACADM. 122 Viewing network interfaces available on host OS using RACADM. 122	DRAC and managed system information	114
Configuring Asset Tracking. 114 Viewing system inventory. 115 Viewing sensor information. 116 Monitoring performance index of CPU, memory, and input output modules. 117 Monitoring performance index of CPU, memory, and input output modules using web interface. 118 Monitoring performance index for of CPU, memory, and input output modules using RACADM. 118 GPU (Accelerators) Management. 112 Checking the system for Fresh Air compliance. 122 Viewing historical temperature data 12 Viewing historical temperature data using iDRAC web interface. 122 Viewing historical temperature data using RACADM. 122 Configuring warning threshold for inlet temperature. 122 Viewing network interfaces available on host OS using RACADM. 122 Viewing network interfaces available on host OS using RACADM. 123 Viewing network interfaces available on host OS using RACADM. 124 Viewing network interfaces available on host OS using RACADM. 123 Viewing network interfaces available on host OS using RACADM. 124 Viewing network interfaces available on host OS using RACADM. 125 Viewing network interfaces available on host OS using RACADM. 12	ystem health and properties	114
Viewing system inventory. 115 Viewing sensor information. 116 Monitoring performance index of CPU, memory, and input output modules. 117 Monitoring performance index of CPU, memory, and input output modules using web interface. 117 Monitoring performance index of CPU, memory, and input output modules using RACADM. 118 Idle Server Detection. 118 GPU (Accelerators) Management. 119 Checking the system for Fresh Air compliance. 120 Viewing historical temperature data 12 Viewing historical temperature data using iDRAC web interface. 122 Viewing historical temperature data using RACADM. 122 Configuring warning threshold for inlet temperature. 122 Viewing network interfaces available on host OS 122 Viewing network interfaces available on host OS using RACADM. 123 Viewing retwork interfaces available on host OS using RACADM. 123 Viewing retwork interfaces available on host OS using RACADM. 123 Viewing retwork interfaces available on host OS using RACADM. 123 Viewing network interfaces available on host OS using RACADM. 124 Viewing network interfaces available on host OS using RACADM. 125 <t< td=""><td>racking</td><td>114</td></t<>	racking	114
Viewing sensor information 116 Monitoring performance index of CPU, memory, and input output modules 117 Monitoring performance index of CPU, memory, and input output modules using web interface 118 Idle Server Detection 116 GPU (Accelerators) Management. 117 Checking the system for Fresh Air compliance. 122 Viewing historical temperature data 122 Viewing historical temperature data using RACADM. 122 Configuring warning threshold for inlet temperature. 122 Viewing network interfaces available on host OS 122 Viewing network interfaces available on host OS using web interface. 122 Viewing network interfaces available on host OS using RACADM. 122 Viewing network interfaces available on host OS using RACADM. 122 Viewing network interfaces available on host OS using RACADM. 123 Viewing retrormating iDRAC sessions. 123 Viewing or terminating iDRAC sessions. 124 Viewing or terminating iDRAC sessions. 125 Viewing ing informating iDRAC communication. 126 Configuring BIOS for serial connection. 126 Configuring BIOS for serial connection. 126 <t< td=""><td>entory</td><td> 115</td></t<>	entory	115
Monitoring performance index of CPU, memory, and input output modules. 11. Monitoring performance index of CPU, memory, and input output modules using web interface. 118. Monitoring performance index for of CPU, memory, and input output modules using RACADM. 118. Idle Server Detection. 116. GPU (Accelerators) Management. 116. Checking the system for Fresh Air compliance. 122. Viewing historical temperature data. 122. Viewing historical temperature data using iDRAC web interface. 122. Viewing historical temperature data using RACADM. 122. Configuring warning threshold for inlet temperature. 122. Viewing network interfaces available on host OS. 122. Viewing network interfaces available on host OS using web interface. 122. Viewing retwork interfaces available on host OS using RACADM. 122. Viewing network interfaces available on host OS using RACADM. 123. Viewing retwork interfaces available on host OS using RACADM. 124. Viewing ing iDRAC sessions. 123. Viewing iDRAC sessions using web interface. 125. Viewing network interfaces available on host OS using RACADM. 125. Configuring iDRAC sessions using web interface.	rmation	116
Monitoring performance index of CPU, memory, and input output modules using web interface. 118 Monitoring performance index for of CPU, memory, and input output modules using RACADM. 118 Idle Server Detection. 118 GPU (Accelerators) Management. 112 Checking the system for Fresh Air compliance. 122 Viewing historical temperature data 12 Viewing historical temperature data using iDRAC web interface. 12 Viewing historical temperature data using RACADM. 122 Configuring warning threshold for inlet temperature. 122 Viewing network interfaces available on host OS. 122 Viewing network interfaces available on host OS using web interface. 122 Viewing network interfaces available on host OS using RACADM. 122 Viewing network interfaces available on host OS using RACADM. 123 Viewing network interfaces available on host OS using RACADM. 123 Viewing or terminating iDRAC sessions. 124 Viewing interfaces mezzanine card fabric connections. 125 Viewing interface sessions using web interface. 125 Communicating with iDRAC through serial connection using DB9 cable. 126 Configuring BIOS for serial connection. 126 <	ance index of CPU, memory, and input output modules	117
Monitoring performance index for of CPU, memory, and input output modules using RACADM. 118 Idle Server Detection. 118 GPU (Accelerators) Management. 118 Checking the system for Fresh Air compliance. 120 Viewing historical temperature data 12 Viewing historical temperature data using iDRAC web interface. 12 Viewing historical temperature data using RACADM. 12 Configuring warning threshold for inlet temperature. 122 Viewing network interfaces available on host OS 122 Viewing network interfaces available on host OS using web interface. 122 Viewing network interfaces available on host OS using RACADM. 122 Viewing network interfaces available on host OS using RACADM. 123 Viewing network interfaces available on host OS using RACADM. 123 Viewing network interfaces available on host OS using RACADM. 123 Viewing or terminating iDRAC sessions. 124 Viewing network interfaces available on host OS using RACADM. 125 Viewing network interfaces available on host OS using RACADM. 125 Viewing BlexAddress mezzanine card fabric connections. 126 Configuring BIDRAC sessions using web interface. 126	prmance index of CPU, memory, and input output modules using web interface	118
Idle Server Detection. 118 GPU (Accelerators) Management. 119 Checking the system for Fresh Air compliance. 120 Viewing historical temperature data 121 Viewing historical temperature data 122 Viewing historical temperature data 122 Viewing historical temperature data using RACADM. 122 Configuring warning threshold for inlet temperature. 122 Viewing network interfaces available on host OS. 122 Viewing network interfaces available on host OS using web interface. 122 Viewing network interfaces available on host OS using RACADM. 123 Viewing retwork interfaces available on host OS using RACADM. 124 Viewing network interfaces available on host OS using RACADM. 125 Viewing or terminating iDRAC sessions. 126 Configuring iDRAC sessions using web interface. 126 Configuring BIOS for serial connection. 126 Configuring BIOS for serial connection. 126 Configuring BIOS for serial connection. 126 Enabling IPMI serial connection basic and terminal modes. 127 Switching between RAC serial and serial console while using DB9 cable. 126 Swit	prmance index for of CPU, memory, and input output modules using RACADM	118
GPU (Accelerators) Management. 119 Checking the system for Fresh Air compliance. 120 Viewing historical temperature data 12 Viewing historical temperature data using iDRAC web interface. 12 Viewing historical temperature data using RACADM. 12 Configuring warning threshold for inlet temperature. 12 Viewing network interfaces available on host OS 12 Viewing network interfaces available on host OS using web interface. 12 Viewing network interfaces available on host OS using RACADM. 12 Viewing network interfaces available on host OS using RACADM. 12 Viewing network interfaces available on host OS using RACADM. 12 Viewing retraining iDRAC sessions. 12 Viewing or terminating iDRAC sessions. 12 Terminating iDRAC sessions using web interface. 12 apter 7: Setting up iDRAC communication. 126 Configuring BIOS for serial connection. 126 Configuring BIOS for serial connection. 126 Enabling RAC serial connection. 126 Enabling IPMI serial connection basic and terminal modes. 127 Switching between RAC serial and serial console while using DB9 cable. 125	on	118
Checking the system for Fresh Air compliance. 120 Viewing historical temperature data. 12 Viewing historical temperature data using iDRAC web interface. 12 Viewing historical temperature data using RACADM. 12 Configuring warning threshold for inlet temperature. 12 Viewing network interfaces available on host OS. 12 Viewing network interfaces available on host OS using web interface. 12 Viewing network interfaces available on host OS using RACADM. 12 Viewing network interfaces available on host OS using RACADM. 12 Viewing network interfaces available on host OS using RACADM. 12 Viewing retwork interfaces available on host OS using RACADM. 12 Viewing or terminating iDRAC sessions. 12 Viewing or terminating iDRAC sessions. 12 Terminating iDRAC sessions using web interface. 12 Sommunicating with iDRAC through serial connection using DB9 cable. 12 Configuring BIOS for serial connection. 12 Enabling RAC serial connection 12 Switching between RAC serial and serial console while using DB9 cable. 12 Switching from serial console to RAC serial. 12 Switching from RAC serial to s) Management	119
Viewing historical temperature data. 12 Viewing historical temperature data using iDRAC web interface. 12 Viewing historical temperature data using RACADM. 12 Configuring warning threshold for inlet temperature. 12 Viewing network interfaces available on host OS. 12 Viewing network interfaces available on host OS using web interface. 12 Viewing network interfaces available on host OS using RACADM. 12 Viewing network interfaces available on host OS using RACADM. 12 Viewing network interfaces available on host OS using RACADM. 12 Viewing network interfaces available on host OS using RACADM. 12 Viewing retwork interfaces available on host OS using RACADM. 12 Viewing retwork interfaces available on host OS using RACADM. 12 Viewing network interfaces available on host OS using RACADM. 12 Viewing retwork interfaces available on host OS using RACADM. 12 Viewing ing REXAddress mezzanine card fabric connections. 12 Viewing iDRAC communication 12 Communicating with iDRAC through serial connection using DB9 cable. 12 Configuring BIOS for serial connection 12 Enabling RAC serial connection 12 <	m for Fresh Air compliance	120
Viewing historical temperature data using iDRAC web interface. 12 Viewing historical temperature data using RACADM. 12 Configuring warning threshold for inlet temperature. 12 Viewing network interfaces available on host OS. 122 Viewing network interfaces available on host OS using web interface. 122 Viewing network interfaces available on host OS using RACADM. 122 Viewing network interfaces available on host OS using RACADM. 123 Viewing retwork interfaces available on host OS using RACADM. 123 Viewing retwork interfaces available on host OS using RACADM. 124 Viewing retwork interfaces available on host OS using RACADM. 125 Viewing or terminating iDRAC sessions. 125 Terminating iDRAC sessions using web interface. 125 apter 7: Setting up iDRAC communication. 126 Configuring BIOS for serial connection. 126 Configuring BIOS for serial connection. 126 Enabling RAC serial connection. 126 Enabling IPMI serial connection basic and terminal modes. 127 Switching from serial console to RAC serial. 128 Switching from serial console to RAC serial. 129 Switching from RAC serial to	emperature data	121
Viewing historical temperature data using RACADM	al temperature data using iDRAC web interface.	121
Configuring warning threshold for inlet temperature. 122 Configuring warning threshold for inlet temperature. 122 Viewing network interfaces available on host OS 122 Viewing network interfaces available on host OS using web interface. 122 Viewing network interfaces available on host OS using RACADM. 123 Viewing network interfaces available on host OS using RACADM. 123 Viewing retwork interfaces available on host OS using RACADM. 123 Viewing retwork interfaces available on host OS using RACADM. 123 Viewing retwork interfaces available on host OS using RACADM. 123 Viewing retwork interfaces available on host OS using RACADM. 123 Viewing retwork interfaces available on host OS using RACADM. 123 Viewing retwork interfaces available on host OS using RACADM. 123 Viewing retwork interfaces available on host OS using RACADM. 123 Viewing retwork interfaces available on host OS using RACADM. 124 Viewing retwork interfaces available on host OS using RACADM. 125 Priminating iDRAC sersions using web interface. 126 Communicating with iDRAC through serial connection using DB9 cable. 126 Enabling RAC serial connection. 126 Enabling	al temperature data using RACADM	121
Viewing network interfaces available on host OS. 122 Viewing network interfaces available on host OS using web interface. 122 Viewing network interfaces available on host OS using RACADM. 123 Viewing network interfaces available on host OS using RACADM. 123 Viewing FlexAddress mezzanine card fabric connections. 123 Viewing or terminating iDRAC sessions. 123 Terminating iDRAC sessions using web interface. 123 apter 7: Setting up iDRAC communication. 126 Communicating with iDRAC through serial connection using DB9 cable. 126 Configuring BIOS for serial connection. 126 Enabling RAC serial connection. 126 Enabling IPMI serial connection basic and terminal modes. 127 Switching between RAC serial and serial console while using DB9 cable. 126 Switching from serial console to RAC serial. 126 Switching from RAC serial to serial console. 126 Switching from RAC serial to serial console. 126 Communicating with iDRAC using IPMI SOL. 128	ning threshold for inlet temperature	122
Viewing network interfaces available on host OS using web interface. 122 Viewing network interfaces available on host OS using RACADM. 123 Viewing FlexAddress mezzanine card fabric connections. 123 Viewing or terminating iDRAC sessions. 123 Terminating iDRAC sessions using web interface. 123 apter 7: Setting up iDRAC communication. 126 Communicating with iDRAC through serial connection using DB9 cable. 126 Configuring BIOS for serial connection. 126 Enabling RAC serial connection basic and terminal modes. 127 Switching between RAC serial and serial console while using DB9 cable. 126 Switching from serial console to RAC serial. 126 Switching from RAC serial to serial console. 127 Switching from RAC serial to serial console. 126 Communicating with iDRAC serial to serial console. 127 Switching from RAC serial to serial console. 126 Switching from RAC serial to serial console. 126 Communicating with iDRAC using IPMI SOL. 128 Communicating with iDRAC using IPMI SOL. 129	terfaces available on host OS	122
Viewing network interfaces available on host OS using RACADM. 122 Viewing FlexAddress mezzanine card fabric connections. 123 Viewing or terminating iDRAC sessions. 123 Terminating iDRAC sessions using web interface. 123 apter 7: Setting up iDRAC communication. 126 Communicating with iDRAC through serial connection using DB9 cable. 126 Configuring BIOS for serial connection. 126 Enabling RAC serial connection basic and terminal modes. 127 Switching between RAC serial and serial console while using DB9 cable. 126 Switching from serial console to RAC serial. 126 Switching from RAC serial to serial console. 126 Communicating with iDRAC using IPMI SOL. 126	<pre>/ interfaces available on host OS using web interface</pre>	122
Viewing FlexAddress mezzanine card fabric connections 123 Viewing or terminating iDRAC sessions 123 Terminating iDRAC sessions using web interface 123 apter 7: Setting up iDRAC communication 126 Communicating with iDRAC through serial connection using DB9 cable 126 Configuring BIOS for serial connection 126 Enabling RAC serial connection 126 Switching between RAC serial and serial console while using DB9 cable 126 Switching from serial console to RAC serial 126 Switching from RAC serial to serial console 126 Communicating with iDRAC using IPMI SOL 128	terfaces available on host OS using RACADM	122
Viewing or terminating iDRAC sessions. 123 Terminating iDRAC sessions using web interface. 123 apter 7: Setting up iDRAC communication. 126 Communicating with iDRAC through serial connection using DB9 cable. 126 Configuring BIOS for serial connection. 126 Enabling RAC serial connection. 126 Enabling IPMI serial connection basic and terminal modes. 127 Switching between RAC serial and serial console while using DB9 cable. 126 Switching from serial console to RAC serial. 126 Switching from RAC serial to serial console. 126 Switching with iDRAC using IPMI SOL. 126 Communicating with iDRAC using IPMI SOL. 126	remarkable of host of dailing (ACADIVIIIIIIIIIIIIIIIIIIIIIIIIIIIIIIIIII	123
Image: Section of the minating iDRAC sessions using web interface. 12 Image: Section of the minating iDRAC sessions using web interface. 12 Image: Section of the minating iDRAC sessions using web interface. 12 Image: Section of the minating iDRAC sessions using web interface. 12 Image: Section of the minating iDRAC communication. 12 Image: Section of the minating iDRAC through serial connection using DB9 cable. 12 Image: Section of the minating iDRAC through serial connection using DB9 cable. 12 Image: Section of the minating iDRAC through serial connection using DB9 cable. 12 Image: Section of the minating iDRAC through serial connection. 12 Image: Section of the minating iDRAC through serial connection. 12 Image: Section of the minating iDRAC through serial console while using DB9 cable. 12 Image: Section of the minating iDRAC serial console to RAC serial. 12 Image: Section of the minating iDRAC using IPMI SOL. 12		107
apter 7: Setting up iDRAC communication	AC sessions using web interface	123
Communicating with iDRAC through serial connection using DB9 cable. 126 Configuring BIOS for serial connection. 126 Enabling RAC serial connection. 126 Enabling IPMI serial connection basic and terminal modes. 127 Switching between RAC serial and serial console while using DB9 cable. 129 Switching from serial console to RAC serial. 129 Switching from RAC serial to serial console. 129 Communicating with iDRAC using IPMI SOL. 129	p iDRAC communication	. 125
Configuring BIOS for serial connection	h iDRAC through serial connection using DR9 cable	126
Enabling RAC serial connection basic and terminal modes	S for serial connection	126
Enabling IPMI serial connection basic and terminal modes		120
Switching from serial console to RAC serial	rial connection basic and terminal modes	107
Switching from serial console to RAC serial console while using DB9 cable	PAC aprial and aprial appeals while using DPA apple	ız/
Switching from RAC serial to serial console	RAC serial and serial console while using DB9 Cable	100
Communicating with iDRAC using IPMI SOL	Serial Console LO RAC Serial	100
Communicating with IDRAC using IPMI SOL129		129
	n idrac using IPMI SOL	129

OTOCO

	IN
Configuring iDRAC to use SOL	130
Enabling supported protocol	131
Communicating with iDRAC using IPMI over LAN	1.34
Configuring IPMI over LAN using web interface	1.3.4
Configuring IPMI over LAN using iDRAC settings utility	134
Configuring IPML over LAN using RACADM	134
Enabling or disabling remote RACADM	135
Enabling or disabiling remote RACADM using web interface	135
Enabling of disabling remote RACADM using Web Interface	135
	135
Enabling IDMI on managed system	
Configuring Linux for serial console during boot in RHEL 6	136
Enabling login to the virtual console after boot	136
Configuring sorial terminal in RHEL 7	
Supported SSH ervetography schemes	130
Light public key authorization for SSH	
	140
Chapter 8: Configuring user accounts and privileges	
iDRAC user roles and privileges	
Recommended characters in user names and passwords	
Configuring local users	
Configuring local users using iDRAC web interface	
Configuring local users using RACADM	
Configuring Active Directory users	
Prerequisites for using Active Directory authentication for iDRAC.	
Supported Active Directory authentication mechanisms	
Standard schema Active Directory overview	148
Configuring Standard schema Active Directory	150
Extended schema Active Directory overview.	
Configuring Extended schema Active Directory	154
Testing Active Directory settings	161
Configuring generic I DAP users	
Configuring generic LDAP directory service using iDRAC web-based interface	162
Configuring generic LDAP directory service using RACADM	162
Testing LDAP directory service settings	
Chapter 9: System Configuration Lockdown mode	
Chapter 10: Configuring iDRAC for Single Sign-On or smart card login	
Prerequisites for Active Directory Single Sign-On or smart card login	
Registering iDRAC on Domain name System	
Creating Active Directory objects and providing privileges	
Configuring iDRAC SSO login for Active Directory users	
Creating a User in Active Directory for SSO	167
Generating Kerberos keytab file	168
Configuring iDRAC SSO login for Active Directory users using web interface	168
Configuring iDRAC SSO login for Active Directory users using RACADM	

FIS.<u>1049</u> Mov. <u>33</u>



Enabling or disabling smart card login	. 169
Enabling or disabling smart card login using web interface	169
Enabling or disabling smart card login using RACADM	169
Enabling or disabling smart card login using iDRAC settings utility	169
Configuring Smart Card Login	170
Configuring iDRAC smart card login for Active Directory users	. 170
Configuring iDRAC smart card login for local users	. 170
Using Smart Card to Login	171

Chapter 11: Configuring iDRAC to send alerts	172
Enabling or disabling alerts	
Enabling or disabling alerts using web interface	172
Enabling or disabling alerts using RACADM	173
Enabling or disabling alerts using iDRAC settings utility	173
Filtering alerts	173
Filtering alerts using iDRAC web interface	
Filtering alerts using RACADM	174
Setting event alerts	174
Setting event alerts using web interface	
Setting event alerts using RACADM	
Setting alert recurrence event	
Setting alert recurrence events using RACADM	174
Setting alert recurrence events using iDRAC web interface	175
Setting event actions	175
Setting event actions using web interface	
Setting event actions using RACADM	
Configuring email alert, SNMP trap, or IPMI trap settings	175
Configuring IP alert destinations	
Configuring email alert settings	
Configuring WS Eventing	179
Configuring Redfish Eventing	
Monitoring chassis events	
Monitoring chassis events using the iDRAC web interface	180
Monitoring chassis events using RACADM	180
Alerts message IDs	

Chapter 12: iDRAC 9 Group Manager	
Group Manager	
Summary View	
Network Configuration requirements	
Manage Logins	
Add a New User	
Change User Password	
Delete User	
Configure Alerts	
Export	
Discovered Servers View	
Jobs View	
Jobs Export	
•	



Group Information Panel	
Group Settings	
Actions on a selected Server	191
iDRAC Group Firmware Update	

Chapter 13: Managing logs	
Viewing System Event Log	
Viewing System Event Log using web interface	
Viewing System Event Log using RACADM	
Viewing System Event Log using iDRAC settings utility	
Viewing Lifecycle log	
Viewing Lifecycle log using web interface	
Viewing Lifecycle log using RACADM	
Exporting Lifecycle Controller logs	
Exporting Lifecycle Controller logs using web interface	
Exporting Lifecycle Controller logs using RACADM	
Adding work notes	
Configuring remote system logging	
Configuring remote system logging using web interface	
Configuring remote system logging using RACADM	

Chapter 14: Monitoring and managing power in iDRAC	197
Monitoring power	197
Monitoring performance index of CPU, memory, and input output modules using web interface	197
Monitoring performance index for of CPU, memory, and input output modules using RACADM	198
Setting warning threshold for power consumption	198
Setting warning threshold for power consumption using web interface	198
Executing power control operations	198
Executing power control operations using web interface	199
Executing power control operations using RACADM	199
Power capping	199
Power capping in Blade servers	199
Viewing and configuring power cap policy	199
Configuring power supply options	200
Configuring power supply options using web interface	201
Configuring power supply options using RACADM	201
Configuring power supply options using iDRAC settings utility	201
Enabling or disabling power button	201
Multi-Vector Cooling	201
Chapter 15: iDRAC Direct Updates	203
Chapter 16: Inventorying, monitoring, and configuring network devices	204
Inventorying and monitoring network devices	204
Monitoring network devices using web interface	204
Monitoring network devices using RACADM	204
Connection View	205
Inventorying and monitoring FC HBA devices	207

Manitaring FOLUDA deviage using DACADM	
Monitoring FC HBA devices using RACADM	
Menitering SED Transceiver devices web interface	207
Monitoring SFP Transceiver devices using web interface	208
Monitoring SFP Transceiver devices using RACADIVI	208
Periel Data Caratura	208
Serial Data Capture	
Dynamic configuration of virtual addresses, initiator, and storage target settings	
Supported cards for IO Identity Optimization	
Supported NIC firmware versions for IO Identity Optimization	
Virtual or Remote assigned Address and Persistence Policy behavior when iDRAC is set to Remote-Assigned Address mode or Console mode	212
System behavior for FlexAddress and IO Identity	213
Enabling or disabling IO Identity Optimization	
SSD Wear Threshold	
Configuring persistence policy settings	215
Chapter 17: Managing storage devices	219
Understanding RAID concepts	220
What is RAID	
Organizing data storage for availability and performance	
Choosing RAID levels	
Comparing RAID level performance	
Supported controllers	229
Supported enclosures	2.30
Summary of supported features for storage devices	230
Inventorving and monitoring storage devices	235
Monitoring storage devices using web interface	236
Monitoring storage devices using RACADM	236
Monitoring backplane using iDRAC settings utility	236
Viewing storage device topology	236
Managing physical disks	
Assigning or uppesigning physical disk on global bot appro-	
Assigning of unassigning physical disk as global not spare	2J7
Converting a physical disk to RAID of non-RAID mode	
El asing physical disks	
El asing SED/ISE device data	
Rebuild Physical Disk.	
Creating virtual disks.	
Greating virtual disks	
Ealung virtual disk cache policies	
Deleting Virtual disks	
Checking virtual disk consistency	
Initializing virtual disks	
Encrypting virtual disks	245
Assigning or unassigning dedicated hot spares	
Managing virtual disks using web interface	247
Managing virtual disks using RACADM	248
RAID Configuration Features	
Managing controllers	250
Configuring controller properties	250
Importing or auto importing foreign configuration	253

OTOCO

10 Contents

	(TE)
Clearing foreign configuration	254
Resetting controller configuration	
Switching the controller mode	
12 Gbps SAS HBA adapter operations	
Monitoring predictive failure analysis on drives	
Controller operations in non-RAID mode or HBA mode	
Running RAID configuration jobs on multiple storage controllers	
Manage Preserved cache	
Managing PCIe SSDs	
Inventorying and monitoring PCIe SSDs	
Preparing to remove PCIe SSD	
Erasing PCIe SSD device data	
Managing enclosures or backplanes	
Configuring backplane mode	
Viewing universal slots	
Setting SGPIO mode	
Set Enclosure Asset Tag	
Set Enclosure Asset Name	
Choosing operation mode to apply settings	
Choosing operation mode using web interface	
Choosing operation mode using RACADM	
Viewing and applying pending operations	
Viewing, applying, or deleting pending operations using web interface	
Viewing and applying pending operations using RACADM	
Storage devices — apply operation scenarios	
Blinking or unblinking component LEDs	
Blinking or unblinking component LEDs using web interface	269
Blinking or unblinking component LEDs using RACADM	
Warm reboot	
Chapter 18: BIOS Settings	271
BIOS Live Scanning	272
BIOS Recovery and Hardware Root of Trust (RoT)	273
Chanter 10, Configuring and using vistual concele	074
Chapter 19: Configuring and using virtual console	2/4 275
	275
Configuring virtual console using web interface	
Configuring virtual console using RACADM	
Freviewing virtual console	
	270 777
	277 777
Disabling warning massages while launahing virtual console or virtual madia using lava a	
bisabiling warning messages while launching virtual console of virtual media using Java o	278
Using virtual console viewer	
eHTML5 based virtual console	
HTML5 based virtual console	281
Synchronizing mouse pointers	
Passing all keystrokes through virtual console for Java or ActiveX plua-in	

FIS.1053



Chapter 20: Using iDRAC Service Module	
Installing iDRAC Service Module	
Installing iDRAC Service Module from iDRAC Express and Basic	
Installing iDRAC Service Module from iDRAC Enterprise	
Supported operating systems for iDRAC Service Module	
iDRAC Service Module monitoring features	
Using iDRAC Service Module from iDRAC web interface	294
Using iDRAC Service Module from RACADM	
Chapter 21: Using USB port for server management	
Accessing iDRAC interface over direct USB connection	
Configuring iDRAC using server configuration profile on USB device	296
Configuring USB management port settings	
Importing Server Configuration Profile from USB device	
Chapter 22: Using Quick Sync 2	300
Configuring iDRAC Quick Sync 2	
Configuring iDRAC Quick Sync 2 settings using web interface	
Configuring iDRAC Quick Sync 2 settings using RACADM	
Configuring iDRAC Quick Sync 2 settings using iDRAC settings utility	
Using mobile device to view iDRAC information	
Chapter 23: Managing virtual media	302
Supported drives and devices	
Configuring virtual media	
Configuring virtual media using iDRAC web interface	
Configuring virtual media using RACADM	
Configuring virtual media using iDRAC settings utility	
Attached media state and system response	
Accessing virtual media	
Launching virtual media using virtual console	
Launching virtual media without using virtual console	
Adding virtual media images	
Viewing virtual device details	
Accessing drivers	
Resetting USB	
Mapping virtual drive	
Unmapping virtual drive	
Setting boot order through BIOS	
Enabling boot once for virtual media	
Chapter 24: Managing vFlash SD card	309
Configuring vFlash SD card	
Viewing vFlash SD card properties	
Enabling or disabling vFlash functionality	
Initializing vFlash SD card	
Getting the last status using RACADM	
Managing vFlash partitions	



Creating an empty partition	
Creating a partition using an image file	
Formatting a partition	
Viewing available partitions	
Modifying a partition	
Attaching or detaching partitions	
Deleting existing partitions	
Downloading partition contents	
Booting to a partition	
Chapter 25: Using SMCLP	
System management capabilities using SMCLP	
Running SMCLP commands	
iDRAC SMCLP syntax	
Navigating the map address space	
Using show verb	
Using the -display option	
Using the -level option	
Using the -output option	
Usage examples	
Server power management	
SEL management	
Map target navigation	
Chapter 26: Deploying operating systems	307
enchaer zei Behiedung ebergrung ederengenen unter einen ein	
Deploying operating system using remote file share	
Deploying operating system using remote file share Managing remote file shares	
Deploying operating system using remote file share Managing remote file shares Configuring remote file share using web interface	
Deploying operating system using remote file share Managing remote file shares Configuring remote file share using web interface Configuring remote file share using RACADM	
Deploying operating system using remote file share Managing remote file shares Configuring remote file share using web interface Configuring remote file share using RACADM Deploying operating system using virtual media	
Deploying operating system using remote file share Managing remote file shares Configuring remote file share using web interface Configuring remote file share using RACADM Deploying operating system using virtual media Installing operating system from multiple disks	
Deploying operating system using remote file share Managing remote file shares Configuring remote file share using web interface Configuring remote file share using RACADM Deploying operating system using virtual media Installing operating system from multiple disks Deploying embedded operating system on SD card	
Deploying operating system using remote file share Managing remote file shares Configuring remote file share using web interface Configuring remote file share using RACADM Deploying operating system using virtual media Installing operating system from multiple disks Deploying embedded operating system on SD card Enabling SD module and redundancy in BIOS	
Deploying operating system using remote file share	
Deploying operating system using remote file share	327 327 328 328 329 329 329 330 330 330 330 330 331
Deploying operating system using remote file share	327 327 328 328 329 329 330 330 330 330 330 330 331 331 331
Deploying operating system using remote file share	
Deploying operating system using remote file share	327
Deploying operating system using remote file share. Managing remote file shares. Configuring remote file share using web interface. Configuring remote file share using RACADM. Deploying operating system using virtual media. Installing operating system from multiple disks. Deploying embedded operating system on SD card. Enabling SD module and redundancy in BIOS. Chapter 27: Troubleshooting managed system using iDRAC. Using diagnostic console. Reset iDRAC and Reset iDRAC to default Scheduling remote automated diagnostics using RACADM. Viewing post codes.	327 327 328 329 329 329 330 330 330 330 330 331 331 331 331 332 332 332 332
Deploying operating system using remote file share	
Deploying operating system using remote file share	
Deploying operating system using remote file share	327 327 328 329 329 329 330 330 330 330 330 331 331 331 332 332 332 332 333 333 333
Deploying operating system using remote file share	
Deploying operating system using remote file share	327
Deploying operating system using remote file share	327 327 328 329 329 329 330 330 330 330 330 331 331 331 332 332 332 333 333 333 333
Deploying operating system using remote file share Managing remote file shares Configuring remote file share using web interface Configuring remote file share using RACADM Deploying operating system using virtual media Installing operating system from multiple disks Deploying embedded operating system on SD card Enabling SD module and redundancy in BIOS Chapter 27: Troubleshooting managed system using iDRAC Using diagnostic console Reset iDRAC and Reset iDRAC to default Scheduling remote automated diagnostics Scheduling remote automated diagnostics using RACADM Viewing post codes Viewing boot and crash capture videos Configuring video capture settings Viewing logs Viewing logs Viewing system front panel LCD status Viewing system front panel LED status	327
Deploying operating system using remote file share Managing remote file shares Configuring remote file share using web interface Configuring remote file share using RACADM Deploying operating system using virtual media Installing operating system from multiple disks. Deploying embedded operating system on SD card Enabling SD module and redundancy in BIOS Chapter 27: Troubleshooting managed system using iDRAC Using diagnostic console Reset iDRAC and Reset iDRAC to default Scheduling remote automated diagnostics using RACADM Viewing post codes Viewing boot and crash capture videos Configuring video capture settings Viewing logs Viewing logs Viewing logs Viewing system front panel LCD status Viewing system front panel LED status Hardware trouble indicators	327 327 328 329 329 329 330 330 330 330 330 330 331 331 331 331



	1
Checking server status screen for error messages	
Restarting iDRAC	
Reset to Custom Defaults (RTD)	
Resetting iDRAC using iDRAC web interface	
Resetting iDRAC using RACADM	
Erasing system and user data	
Resetting iDRAC to factory default settings	
Resetting iDRAC to factory default settings using iDRAC web interface	
Resetting iDRAC to factory default settings using iDRAC settings utility	
Chapter 28: SupportAssist Integration in iDRAC	339
SupportAssist Registration	
Installing Service Module	
Server OS Proxy Information	
SupportAssist	
Service Request Portal	
Collection Log	
Generating SupportAssist Collection	
Generating SupportAssist Collection manually using iDRAC web interface	
Settings	
Collection Settings	
Contact Information	
Chanter 20: Frequently solved questions	7 4 7
Chapter 29: Frequently asked questions	
System Event Log	

System Event Log	343
Custom sender email configuration for iDRAC alerts	
Network security	344
Telemetry streaming	
Active Directory	
Single Sign-On	
Smart card login	
Virtual console	
Virtual media	
vFlash SD card	352
SNMP authentication	352
Storage devices	352
GPU (Accelerators)	353
iDRAC Service Module	353
RACADM	
Permanently setting the default password to calvin	355
Miscellaneous	356

Chapter 30: Use case scenarios	361
- Troubleshooting an inaccessible managed system	
Obtaining system information and assess system health	
Setting up alerts and configuring email alerts	
Viewing and exporting System Event Log and Lifecycle Log	
Interfaces to update iDRAC firmware	
Performing graceful shutdown	



Creating new administrator user account	.362
Launching servers remote console and mounting a USB drive	363
Installing bare metal OS using attached virtual media and remote file share	363
Managing rack density	.363
Installing new electronic license	363
Applying IO Identity configuration settings for multiple network cards in single host system reboot	363



Overview of iDRAC

The Integrated Dell Remote Access Controller (iDRAC) is designed to make you more productive as a system administrator and improve the overall availability of Dell EMC servers. iDRAC alerts you to system issues, helps you to perform remote management, and reduces the need for physical access to the system.

iDRAC technology is part of a larger data center solution that increases availability of business critical applications and workloads. The technology allows you to deploy, monitor, manage, configure, update, and troubleshoot Dell EMC systems from any location without using any agents or an operating system.

Several products work with the iDRAC to simplify and streamline IT operations. Following are some of the tools:

- OpenManage Enterprise
- OpenManage Power Center Plug in
- OpenManage Integration for VMware vCenter
- Dell Repository Manager

iDRAC is available in the following variants:

- iDRAC Basic Available by default for 100-500 series servers
- iDRAC Express Available by default on all 600 and higher series of rack or tower servers, and all blade servers
- iDRAC Enterprise Available on all server models
- iDRAC Datacenter Available on all server models

Topics:

- Benefits of using iDRAC
- Key features
- New features added
- How to use this guide
- Supported web browsers
- iDRAC licenses
- Licensed features in iDRAC9
- Interfaces and protocols to access iDRAC
- iDRAC port information
- Other documents you may need
- Contacting Dell
- Accessing documents from Dell support site
- Accessing Redfish API Guide

Benefits of using iDRAC

The benefits include:

- Increased Availability Early notification of potential or actual failures that help prevent a server failure or reduce recovery time after failure.
- Improved Productivity and Lower Total Cost of Ownership (TCO) Extending the reach of administrators to larger numbers of distant servers can make IT staff more productive while driving down operational costs such as travel.
- Secure Environment By providing secure access to remote servers, administrators can perform critical management functions while maintaining server and network security.
- Enhanced Embedded Management through Lifecycle Controller Lifecycle Controller provides deployment and simplified serviceability through Lifecycle Controller GUI for local deployment and Remote Services (WSMan) interfaces for remote deployment integrated with Dell OpenManage Enterprise and partner consoles.

For more information about Lifecycle Controller GUI, see *Lifecycle Controller User's Guide* and for remote services, see *Lifecycle Controller Remote Services Quick Start Guide* available at https://www.dell.com/idracmanuals.



Key features

The key features of iDRAC include:

NOTE: Some features are available only with iDRAC Enterprise or Datacenter license. For information on the features available for a license, see iDRAC licenses .

Inventory and Monitoring

- Telemetry data streaming.
- View managed server health.
- Inventory and monitor network adapters and storage subsystem (PERC and direct attached storage) without any operating system agents.
- View and export system inventory.
- View sensor information such as temperature, voltage, and intrusion.
- Monitor CPU state, processor automatic throttling, and predictive failure.
- View memory information.
- Monitor and control power usage.
- Support for SNMPv3 gets and alerts.
- For blade servers: launch Management Module web interface, view OpenManage Enterprise (OME) Modular information, and WWN/MAC addresses.

NOTE: CMC provides access to iDRAC through the M1000E Chassis LCD panel and local console connections. For more information, see *Chassis Management Controller User's Guide* available at https://www.dell.com/cmcmanuals.

- View network interfaces available on host operating systems.
- iDRAC9 provides improved monitoring and management functionality with Quick Sync 2. You need OpenManage Mobile app configured in your Android or iOS mobile device.

Deployment

- Manage vFlash SD card partitions.
- Configure front panel display settings.
- Manage iDRAC network settings.
- Configure and use virtual console and virtual media.
- Deploy operating systems using remote file share, and virtual media.
- Enable auto-discovery.
- Perform server configuration using the export or import XML or JSON profile feature through RACADM, WSMan
- and Redfish. For more information, see *Lifecycle Controller Remote Services Quick Start Guide* available at https://www.dell.com/idracmanuals.
- Configure persistence policy for virtual addresses, initiator, and storage targets.
- Remotely configure storage devices attached to the system at run-time.
- Perform the following operations for storage devices:
 - Physical disks: Assign or unassign physical disk as a global hot spare.
 - Virtual disks:
 - Create virtual disks.
 - Edit virtual disks cache policies.
 - Check virtual disk consistency.
 - Initialize virtual disks.
 - Encrypt virtual disks.
 - Assign or unassign dedicated hot spare.
 - Delete virtual disks.
 - Controllers:
 - Configure controller properties.
 - Import or auto-import foreign configuration.
 - Clear foreign configuration.
 - Reset controller configuration.
 - Create or change security keys.
 - PCIe SSD devices:
 - Inventory and remotely monitor the health of PCIe SSD devices in the server.
 - Prepare the PCIe SSD to be removed.
 - Securely erase the data.



- Set the backplane mode (unified or split mode).
- Blink or unblink component LEDs.
- Apply the device settings immediately, at next system reboot, at a scheduled time, or as a pending operation to be applied as a batch as part of the single job.

Update

- Manage iDRAC licenses.
- Update BIOS and device firmware for devices supported by Lifecycle Controller.
- Update or rollback iDRAC firmware and Lifecycle Controller firmware using a single firmware image.
- Manage staged updates.
- Access iDRAC interface over direct USB connection.
- Configure iDRAC using Server Configuration Profiles on USB device.

Maintenance and Troubleshooting

- Perform power-related operations and monitor power consumption.
- Optimize system performance and power consumption by modifying the thermal settings.
- No dependency on OpenManage Server Administrator for generation of alerts.
- Log event data: Lifecycle and RAC logs.
- Set email alerts, IPMI alerts, remote system logs, WS Eventing logs, Redfish event, and SNMP traps (v1, v2c, and v3) for events and improved email alert notification.
- Capture last system crash image.
- View boot and crash capture videos.
- Out-of-band monitor and alert the performance index of CPU, memory, and I/O modules.
- Configure warning threshold for inlet temperature and power consumption.
- Use iDRAC Service Module to:
 - View operating system information.
 - Replicate Lifecycle Controller logs to operating system logs.
 - Automate system recovery options.
 - Enable or disable status of Full Power Cycle for all System components except the PSU.
 - Remotely hard-reset iDRAC
 - Enable in-band iDRAC SNMP alerts
 - Access iDRAC using host OS (experimental feature)
 - Populate Windows Management Instrumentation (WMI) information.
 - Integrate with SupportAssist collection. This is applicable only if iDRAC Service Module Version 2.0 or later is installed.
 - Generate SupportAssist collection in the following ways:
 - Automatic Using iDRAC Service Module that automatically invokes the OS Collector tool.

Dell Best Practices regarding iDRAC

- Dell iDRAC's are intended to be on a separate management network; they are not designed nor intended to be placed on or connected directly to the Internet. Doing so could expose the connected system to security and other risks for which Dell is not responsible.
- Dell EMC recommends using the Dedicated Gigabit Ethernet port available on rack and tower servers. This interface is not shared with the host operating system and routes the management traffic to a separate physical network, enabling it to be separated from the application traffic. This option implies that iDRAC's dedicated network port routes its traffic separately from the server's LOM or NIC ports. The Dedicated option allows iDRAC to be assigned an IP address from the same subnet or different subnet in comparison to the IP addresses assigned to the Host LOM or NICs.
- Along with locating iDRACs on a separate management subnet, users should isolate the management subnet/vLAN with technologies such as firewalls, and limit access to the subnet/vLAN to authorized server administrators.

Secure Connectivity

Securing access to critical network resources is a priority. iDRAC implements a range of security features that includes:

- Custom signing certificate for Secure Socket Layer (SSL) certificate.
- Signed firmware updates.
- User authentication through Microsoft Active Directory, generic Lightweight Directory Access Protocol (LDAP) Directory Service, or locally administered user IDs and passwords.
- Two-factor authentication using the Smart–Card logon feature. The two-factor authentication is based on the physical smart card and the smart card PIN.
- Single Sign-On and Public Key Authentication.
- Role-based authorization, to configure specific privileges for each user.



- SNMPv3 authentication for user accounts stored locally in the iDRAC. It is recommended to use this, but it is disabled by default.
- User ID and password configuration.
- Default login password modification.
- Set user passwords and BIOS passwords using one-way hash format for improved security.
- FIPS 140-2 Level 1 capability.
- Session time-out configuration (in seconds).
- Configurable IP ports (for HTTP, HTTPS, SSH, Virtual Console, and Virtual Media).
- Secure Shell (SSH) that uses an encrypted transport layer for higher security.
- Login failure limits per IP address, with login blocking from that IP address when the limit is exceeded.
- Limited IP address range for clients connecting to iDRAC.
- Dedicated Gigabit Ethernet adapter available on rack and tower servers (additional hardware may be required).

New features added

This section provides the list of new features added in the following releases:

Firmware version 4.40.00.00

This release includes all the features from the previous releases. Following are the new features that are added in this release:

NOTE: For information about supported systems, refer to the respective version of Release Notes available at https://www.dell.com/support/article/sln308699.

- Added support for enhanced HTML5 (eHTML5) virtual KVM feature in virtual console
- Added support for eHTML5 virtual media
- Enhancement in Storage GUI page
- Added support for direct updates SEP backplane
- Added support for new update for PSU update
- Added support for uploading custom defaults and reset iDRAC to default settings using custom defaults
- Enhanced system lockdown mode support for supported devices

Following are the list of other features added in this release:

- Automation
- Support for Redfish Updates
- Monitoring/Alerting/Troubleshooting
 - FPGA Monitoring
 - SMART data logs enhancements including historical recording
 - Discrete voltage sensor reporting
 - Report actual start and completion info for job queue entries which require a server reboot to apply (Example: BIOS update).
 - Providing CPU serial numbers in SupportAssist Collection
 - **Telemetry** (requires iDRAC Datacenter license)
 - Multi-client support
 - Granular metric report options
 - Provision to POST a new custom MRD (Metric Report Definition) using any of the available 193 Metric Definitions and set desired Report Interval (referred as Recurrence Interval in MRD)
 - A single MRD can have a maximum of 68 Metric Definitions (Metric IDs)
 - Provision to create up to 24 new custom MRDs which in turn will have 24 new Metric Reports. An iDRAC can support a maximum of 48 Metric Reports (24 Pre-canned and 24 Custom
- Security
 - Automatic Certificate Enrollment Enhancements (requires iDRAC Datacenter License)
 - Integrate RSA SecurID Client into iDRAC for 2FA (requires iDRAC Datacenter License)
 - Compliance with STIG requirement "network device must authenticate NTP"
 - Removal of Telnet and TLS 1.0 from web server

• Platform feature support

BOSS 1.5 updates



Infiniband support

In 4.40.00.00 release, following features are added in Storage page on iDRAC GUI:

- From the Dashboard, you can see suggested actions to solve any health alters.
- The Storage page has been modified to included tabs for storage monitoring information, a Storage Hardware and Software Inventory, a list of Pending and Current storage jobs, and SEKM.
 - From the Storage Inventory, users can find all storage related hardware and software.
 - \circ The Pending and Current Jobs tab allows users to queue and monitor jobs from a centralized location.
 - You can also configure SEKM via the Storage page.
- When monitoring storage devices, you can customize the columns that are displayed for each device table. Column customization will be saved and persist between user sessions.
- New basic and advanced filters provided on each device page allow you to easily and efficiently customize the list of objects displayed.
- The Storage Configuration wizard has two options to create a Virtual disk Basic and Advanced.
 - In the basic Virtual Disk wizard, you can quickly create a VD from a list of available RAID configurations. iDRAC will automatically set the default values of the VD to streamline the process.
 - For the Advanced Virtual Disk wizard, you can select all the details of the VD. You can create a new volume for the VD or select and existing volume.
- Each devices pages has new global actions that allow you to show related devices or preform group operations.
 - For example, you can choose physical disks and perform group operation such as Blink, Unblink, and Create Virtual Disk.
 - Also, you can view the Physical disk inventory and create a Virtual Disk by choosing the drives without having to navigate away from the screen.
- Instead of the numerical value, the size of the physical disk is shown as a data visualization with values on the scale.
- This gives you an idea of used and available space on the drive.
- You can filter disks based on the various physical disk properties.
 - The filtering properties are displayed so that the user knows what filtering is currently being applied.

Firmware version 4.30.30.30

This release includes all the features from the previous releases. Following are the new features that are added in this release:

() NOTE: For information about supported systems, refer to the respective version of Release Notes available at https://www.dell.com/support/article/sln308699.

- Added support for PERC 11 for AMD systems
- Added support for NVMe drives behind PERC 11
- Added support for HBA11 for AMD systems
- Added support for CUPS for AMD systems
- Added support for Boot Optimized storage solution 1.5 (BOSS1.5/BOSS-S2)
- Added support for BOSS 1.5 secure firmware update
- Added support for new Matrox video driver
- Added support for NVMe Opal SEDs
- Added support for HW chain of Trust Secure Boot
- Added InifiniBand Adaptor support for Mellanox CX6
- Added support for 24x NVMe backplane for PowerEdge C6525
- Added support for new Matrox video driver
- Added support for Starlord (ConnectX-6 Dx 100GbE) to iDRAC
- Added FQDD related changes for BOSS-S2/PERC 11/HBA 11
- Added support for storage devices (like M.2 and U.2 but not limited) without a backplane
- Added support for Secure Enterprise Key Management (SEKM) for NVMe drives
- Enhanced iDRAC memory from 512MB to 1024MB
- RESTUI changes for failed Email Sending when authentication is disabled


Firmware version 4.20.20.20

Following features were added in this release:

Power Supply Unit (PSU)

- Support for 1100W ~48W DC PSU.
- Removed 4S PSU restriction.

NICs

- Support (4x 10/25 SFP28) OCP 3.0 Dell part # JTK7F Broadcom.
- Support (4x10/25) MX Mezz, Dell part # DCWFP Broadcom and MX 25G Quad port on MX platform.
- Support for adding Broadcom 10GbE NIC card support to R340.

Accelerators and CPU's

- Support for 2 new GPU cards to the Precision 7920 Rack (Navi10DT/W5700, Navi14DT/W5500).
- Support for Nvidia V100S for PowerEdge.
- Support for new Intel Processors: 6250 and 6256.

NVMe

• Support for Samsung PM 1735 and PM 1733 NVMe PCIe storage.

Automation/Scripting/Telemetry

- Support for Redfish 2018R3, 2019R1 & 2019R2 Features.
- Support of CLI method to retrieve POST codes.
- Support for Report Interval limit on Telemetry CUPS to be increased from 1 minute to 1 hour in Power Manager Plug-in.
- Support for Telemetry (Metric Report Enable/ Disable).
- Support for enhanced user logging with SSH.
- Support for adding tier specification flag to PCI Add IPMI command.

Miscellaneous

- Support for Thermal Sensor Board Cable Detection when C6420 chassis with one or more sleds is powered up.
- Support for displaying Slot number in the SLED GUI for 6420.
- Support for always arm AEP and BPS memory for an ADR flow on AC loss or Global Reset.
- Support for 10x2.5" BP/chassis part number change.
- Support for enabling "Unsupported Config" to the SEL log.

Firmware version 4.10.10.10

Following features are added in this release:

Features supported with Default license

• BIOS recovery and Root of Trust (RoT)



Features supported with Enterprise license

• Secure Enterprise Key Management (SEKM) — Added support for Vormetric Data Security Manager.

Features supported with Datacenter license

• BIOS live scanning — Only for AMD systems.

Firmware version 4.00.00.00

This release includes all the features from the previous releases. Following are the new features that are added in this release:

NOTE: For information about supported systems, refer to the respective version of Release Notes available at https://www.dell.com/support/article/sln308699.

Features supported with Datacenter license

- Telemetry streaming metric reports streamed to an analytics tool
- GPU inventory and monitoring
- Thermal Manage Advanced Power and cooling features
- Auto Certificate Enrollment and renewal for SSL certificates
- Virtual Clipboard Support Cut and Paste of text strings into the remote virtual console desktop
- SFP Transceiver Input/Output monitoring
- SMART logs Storage drives
- System Serial Data Buffer Capture
- Idle Server detection

Features supported with Enterprise or Datacenter license

- Multi Factor Authentication via email
- Agent Free Crash Video Capture (Windows only)
- Connection View for LLDP transmit
- System Lockdown mode new icon in header available from any page
- Group Manager 250 node support
- Enhanced support for Secure Enterprise Key Management (SEKM)

Features supported with default license (iDRAC Basic or iDRAC Express)

GUI Enhancement

- \circ $\;$ Task Summary section in the dashboard $\;$
- Search box in the header
- SupportAssist Collection Viewer displays the output in iDRAC GUI
- API, CLI, and SCP
 - Operating system deployment by Server Configuration Profile (SCP)
 - Enable and disable boot order cotrol to SCP and RACADM
 - New schemas to Redfish APIs
 - Option to change boot source state in SCP
 - Automation for Command/attribute auto completion in RACADM

Alerts and Monitoring

- Custom Sender Email Address for email alerts in SMTP configuration
- Cloud based email server in SMTP
- o SMARTlogs in SupportAssist log collection for hard drives and PCle SSD devices
- Include Part Number of failed component in alert messages
- Security
- 22 Overview of iDRAC



- Multiple IP filtering ranges using RACADM commands only
- iDRAC user password maximum length extended to 40 characters
- SSH Public Keys through SCP
- \circ $\;$ Customizable Security banner to SSH login $\;$
- Force Change Password (FCP) for login

• Storage and Storage Controllers

• Enable PERC to switch to SEKM encryption mode

How to use this guide

The contents of this user's guide enable you to perform various tasks using:

- iDRAC web interface Only the task-related information is provided here. For information about the fields and options, see the *iDRAC Online Help* that you can access from the web interface.
- RACADM The RACADM command or the object that you must use is provided here. For more information, see the *iDRAC* RACADM CLI Guide available at https://www.dell.com/idracmanuals.
- iDRAC Settings Utility Only the task-related information is provided here. For information about the fields and options, see the *iDRAC Settings Utility Online Help* that you can access when you click **Help** in the iDRAC Settings GUI (press <F2> during boot, and then click **iDRAC Settings** on the **System Setup Main Menu** page).
- Redfish Only the task-related information is provided here. For information about the fields and options, see the *iDRAC Redfish API Guide* available at www.api-marketplace.com.

Supported web browsers

iDRAC is supported on the following browsers:

- Internet Explorer/Edge
- Mozilla Firefox
- Google Chrome
- Safari

For the list of supported versions, see the *iDRAC Release Notes* available at https://www.dell.com/idracmanuals.

Supported OS and Hypervisors

iDRAC is supported on the following OS, Hypervisors:

- Microsoft Windows Server and Windows PE
- VMware ESXI
- RedHat Enterprise Linux
- SuSe Linux Enterprise Server

(i) NOTE: For the list of supported versions, see the *iDRAC Release Notes* available at https://www.dell.com/idracmanuals.

iDRAC licenses

iDRAC features are available based on the type of the license. Depending on the system model, iDRAC Basic or iDRAC Express license is installed by default. iDRAC Enterprise license, iDRAC Datacenter license, and iDRAC Secure Enterprise Key Manager (SEKM) license are available as an upgrade and can be purchased anytime. Only licensed features are available in the interfaces that enable you to configure or use iDRAC. For more information, see Licensed features in iDRAC9.

Types of licenses

iDRAC Basic or iDRAC Express are the standard licenses available by default on your system. iDRAC Enterprise and Datacenter licenses includes all the licensed features and can be purchased at any time. The types of upsell offered are:

• 30-day evaluation—Evaluation licenses are duration-based and the timer runs when power is applied to the system. This license cannot be extended.



- Perpetual—The license is bound to the Service Tag and is permanent.
- Following table lists the default license available on the following systems:

iDRAC Basic License	iDRAC Express License	iDRAC Enterprise License	iDRAC Datacenter License
PowerEdge Rack/Tower severs series 100-500	 PowerEdge C41XX PowerEdge FC6XX PowerEdge R6XX PowerEdge R64XX PowerEdge R74XX PowerEdge R74XX4 PowerEdge R74XX PowerEdge R8XX PowerEdge R9XX PowerEdge R9XX PowerEdge R9XX PowerEdge T6XX Dell Precision Rack R7920 	All platforms, with upgrade option	All platforms, with upgrade option

Table 1. Default License

iDRAC Express License	iDRAC Enterprise License	iDRAC Datacenter License
PowerEdge C41XX	All platforms, with upgrade option	All platforms, with upgrade option
PowerEdge FC6XX		
PowerEdge R6XX		
PowerEdge R64XX		
PowerEdge R7XX		
PowerEdge R74XXd		
PowerEdge R74XX		
PowerEdge R8XX		
PowerEdge R9XX		
PowerEdge R9XX		
PowerEdge T6XX		
Dell Precision Rack R7920		

NOTE: The default license available with PowerEdge C64XX systems is BMC. The BMC license was custom made for C64XX systems.

(i) NOTE: Express for Blades license is the default license for PowerEdge M6XX and MXXXX systems.

Methods for acquiring licenses

Use any of the following methods to acquire the licenses:

• Dell Digital Locker — Dell Digital Locker allows you to view and manage your products, software, and licensing information in one location. A link to the Dell Digital Locker is available in DRAC web interface- go to **Configuration** > **Licenses**.

(i) NOTE: To know more about Dell Digital Locker, refer to FAQ on the website.

- Email License is attached to an email that is sent after requesting it from the technical support center.
- Point-of-sale License is acquired while placing the order for a system.

(i) NOTE: To manage licenses or purchase new licenses, go to the Dell Digital Locker.

Acquiring license key from Dell Digital Locker

To obtain the license key from your account, you must first register your product using the registration code that is sent in the order confirmation email. This code must be entered in the **Product Registration** tab after logging into Dell Digital Locker.

From the left pane, click the **Products** or **Order History** tab to view the list of your products. Subscription-based products are listed under **Billing accounts** tab.

To download the license key from your Dell Digital Locker account:

- 1. Sign in to your Dell Digital Locker account.
- 2. From the left pane, click Products.
- 3. Click the product that you want to view.
- 4. Click the product name.
- 5. On the Product management page, click Get Key.
- 6. Follow the instructions on the screen to obtain the license key.

NOTE: If you do not have a Dell Digital Locker account, create an account using the email address provided during your purchase.

NOTE: To generate multiple license keys for new purchases, follow the instructions under Tools > License Activation > Unactivated licenses

License operations

Before you perform the license management tasks, ensure that you acquire the licenses. For more information, see the Methods for acquiring licenses.

(i) NOTE: If you have purchased a system with all the licenses pre-installed, then license management is not required.

You can perform the following licensing operations using iDRAC, RACADM, WSMan, Redfish and Lifecycle Controller-Remote Services for one-to-one license management, and Dell License Manager for one-to-many license management:

- View View the current license information.
- Import After acquiring the license, store the license in a local storage and import it into iDRAC using one of the supported interfaces. The license is imported if it passes the validation checks.

() NOTE: Although you can export the factory-installed license, you cannot import it. To import the license, download the equivalent license from the Digital Locker or retrieve it from the email you received when you purchased the license.

(i) **NOTE:** After importing the license, you need to re-login to iDRAC. This is applicable only for iDRAC web interface.

- Export Exports the installed license. For more information, see the *iDRAC Online Help*.
- Delete Deletes the license. For more information, see the *iDRAC Online Help*.
- Learn More Learn more about an installed license, or the licenses available for a component installed in the server.

NOTE: For the Learn More option to display the correct page, ensure that ***.dell.com** is added to the list of Trusted Sites in the Security Settings. For more information, see the Internet Explorer help documentation.

For one-to-many license deployment, you can use Dell License Manager. For more information, see *Dell License Manager User's Guide* available at https://www.dell.com/esmmanuals.

Following are the user privilege requirements for different license operation:

- Licence View and Export: Login privilege.
- License Import and Delete: Login + Configure iDRAC + Server Control privilege.

Managing licenses using iDRAC web interface

To manage the licenses using the iDRAC web interface, go to **Configuration** > **Licenses**.

The **Licensing** page displays the licenses that are associated to devices, or the licenses that are installed but the device is not present in the system. For more information on importing, exporting, or deleting a license, see the *iDRAC Online Help*.

Managing licenses using RACADM

To manage licenses using RACADM, use the license subcommand. For more information, see the

iDRAC RACADM CLI Guide available at https://www.dell.com/idracmanuals.

Fls.<u>1067</u> Mov. 33



Licensed features in iDRAC9

The following table lists iDRAC9 features that are enabled based on the license purchased:

Table 2. Licensed features in iDRAC9

Feature	iDRAC iDRAC9 iDRAC9 Express iDRAC9 Enterprise 9 Basic Express for Blades		iDRAC9 Enterprise	iDRAC9 Datacenter	
Interfaces / Standards					
iDRAC RESTful API and Redfish	Yes	Yes	Yes	Yes	Yes
IPMI 2.0	Yes	Yes	Yes	Yes	Yes
DCMI 1.5	Yes	Yes	Yes	Yes	Yes
Web-based GUI	Yes	Yes	Yes	Yes	Yes
RACADM command line (local/remote)	Yes	Yes	Yes	Yes	Yes
SSH	Yes	Yes	Yes	Yes	Yes
Serial Redirection	Yes	Yes	Yes	Yes	Yes
WSMan	Yes	Yes	Yes	Yes	Yes
Network Time Protocol	No	Yes	Yes	Yes	Yes
Connectivity		-	-		-
Shared NIC (LOM)	Yes	Yes	N/A	Yes	Yes
Dedicated NIC	Yes	Yes	Yes	Yes	Yes
VLAN tagging	Yes	Yes	Yes	Yes	Yes
IPv4	Yes	Yes	Yes	Yes	Yes
IPv6	Yes	Yes	Yes	Yes	Yes
DHCP	Yes	Yes	Yes	Yes	Yes
DHCP with zero touch	No	No	No	Yes	Yes
Dynamic DNS	Yes	Yes	Yes	Yes	Yes
OS pass-through	Yes	Yes	Yes	Yes	Yes
iDRAC Direct -Front panel USB	Yes	Yes	Yes	Yes	Yes
Connection View	Yes	Yes	No	Yes	Yes
Security		-	-		
Role-based authority	Yes	Yes	Yes	Yes	Yes
Local users	Yes	Yes	Yes	Yes	Yes
SSL encryption	Yes	Yes	Yes	Yes	Yes
Secure Enterprise Key Manager	No	No	No	Yes (with SEKM license)	Yes (with SEKM license)
IP blocking	No	Yes	Yes	Yes	Yes
Directory services (AD, LDAP)	No	No	No	Yes	Yes

Inserido ao protocolo **22.951.206-4** por: **Pedro Henrique Golin Linhares** em: 31/03/2025 12:49. A autenticidade deste documento pode ser validada no endereço: https://www.eprotocolo.pr.gov.br/spiweb/validarDocumento com o código: fc1c44dfe0686d5547ffd691199be699.



Feature	iDRAC 9 Basic	iDRAC9 Express	iDRAC9 Express for Blades	iDRAC9 Enterprise	iDRAC9 Datacenter
Two-factor authentication (smart card)	No	No	No	Yes	Yes
Single sign-On	No	No	No	Yes	Yes
PK authentication (for SSH)	No	Yes	Yes	Yes	Yes
OAuth integration with Web based Authentication services	No	No	No	No	Yes
OpenID Connect for Dell EMC Consoles	No	No	No	No	Yes
FIPS 140-2	Yes	Yes	Yes	Yes	Yes
Secure UEFI boot - certificate management	Yes	Yes	Yes	Yes	Yes
Lock down mode	No	No	No	Yes	Yes
Unique iDRAC default password	Yes	Yes	Yes	Yes	Yes
Customizable Security Policy Banner - login page	Yes	Yes	Yes	Yes	Yes
Easy Multi Factor Authentication	No	No	No	No	Yes
Auto Certificate Enrollment (SSL Certs)	No	No	No	No	Yes
iDRAC Quick Sync 2 - optional auth for read operations	Yes	Yes	Yes	Yes	Yes
iDRAC Quick Sync 2 - add mobile device number to LCL	Yes	Yes	Yes	Yes	Yes
System Erase of internal storage devices	Yes	Yes	Yes	Yes	Yes
Remote Presence			•		
Power control	Yes	Yes	Yes	Yes	Yes
Boot control	Yes	Yes	Yes	Yes	Yes
Serial-over-LAN	Yes	Yes	Yes	Yes	Yes
Virtual Media	No	No	Yes	Yes	Yes
Virtual Folders	No	No	No	Yes	Yes
Remote File Share	No	No	No	Yes	Yes
HTML5 access to Virtual Console	No	No	Yes	Yes	Yes
Virtual Console	No	No	Yes	Yes	Yes
VNC connection to OS	No	No	No	Yes	Yes
Quality/bandwidth control	No	No	No	Yes	Yes



Feature	iDRAC 9 Basic	iDRAC9 Express	iDRAC9 Express for Blades	iDRAC9 Enterprise	iDRAC9 Datacenter	
Virtual Console collaboration (up to six simultaneous users)	No	No	No (One user only)	Yes	Yes	
Virtual Console chat	No	No	No	Yes	Yes	
Virtual Flash partitions	No	No	No	Yes	Yes	
(i) NOTE: vFlash is not av	vailable in i	iDRAC9 for	PowerEdge Rx5xx/Cx	ōxx.		
Group Manager	No	No	No	Yes	Yes	
HTTP / HTTPS support along with NFS/CIFS	Yes	Yes	Yes	Yes	Yes	
Power and Thermal			•	1		
Real-time power meter	Yes	Yes	Yes	Yes	Yes	
Power thresholds and alerts	No	Yes	Yes	Yes	Yes	
Real-time power graphing	No	Yes	Yes	Yes	Yes	
Historical power counters	No	Yes	Yes	Yes	Yes	
Power capping	No	No	No	Yes	Yes	
Power Center integration	No	No	No	Yes	Yes	
Temperature monitoring	Yes	Yes	Yes	Yes	Yes	
Temperature graphing	No	Yes	Yes	Yes	Yes	
PCIe airflow customization (LFM)	No	No	No	No	Yes	
Custom Exhaust Control	No	No	No	No	Yes	
Custom Delta-T control	No	No	No	No	Yes	
System Airflow Consumption	No	No	No	No	Yes	
Custom PCle inlet temperature	No	No	No	No	Yes	
Health Monitoring						
Full agent-free monitoring	Yes	Yes	Yes	Yes	Yes	
Predictive failure monitoring	Yes	Yes	Yes	Yes	Yes	
SNMPv1, v2, and v3 (traps and gets)	Yes	Yes	Yes	Yes	Yes	
Email Alerting	No	Yes	Yes	Yes	Yes	
Configurable thresholds	Yes	Yes	Yes	Yes	Yes	
Fan monitoring	Yes	Yes	Yes	Yes	Yes	

28 Overview of iDRAC



Feature	iDRAC 9 Basic	iDRAC9 Express	iDRAC9 Express for Blades	iDRAC9 Enterprise	iDRAC9 Datacenter
Power Supply monitoring	Yes	Yes	Yes	Yes	Yes
Memory monitoring	Yes	Yes	Yes	Yes	Yes
CPU monitoring	Yes	Yes	Yes	Yes	Yes
RAID monitoring	Yes	Yes	Yes	Yes	Yes
NIC monitoring	Yes	Yes	Yes	Yes	Yes
Optic Inventory	Yes	Yes	Yes	Yes	Yes
Optic Statistics	No	No	No	No	Yes
HD monitoring (enclosure)	Yes	Yes	Yes	Yes	Yes
Out of Band Performance Monitoring	No	No	No	Yes	Yes
Alerts for excessive SSD wear	Yes	Yes	Yes	Yes	Yes
Customizable settings for Exhaust Temperature	Yes	Yes	Yes	Yes	Yes
Serial Console Logs	No	No	No	No	Yes
SMART logs for Storage Drives	No	No	No	No	Yes
Idle Server detection	No	No	No	No	Yes
Telemetry Streaming	No	No	No	No	Yes
(i) NOTE: The OpenMana iDRAC.	ige Enterp	rise Advanc	ed license and the Pow	verManage Plugin support telemet	ry data pulls from the
Update					
Remote agent-free update	Yes	Yes	Yes	Yes	Yes
Embedded update tools	Yes	Yes	Yes	Yes	Yes
Update from repository (Auto-Update)	No	No	No	Yes	Yes
Schedule update from repository	No	No	No	Yes	Yes
Improved PSU firmware updates	Yes	Yes	Yes	Yes	Yes
Deployment and Configu	iration				
Local configuration via F10	Yes	Yes	Yes	Yes	Yes



Feature	iDRAC 9 Basic	iDRAC9 Express	iDRAC9 Express for Blades	iDRAC9 Enterprise	iDRAC9 Datacenter
Embedded OS deployment tools	Yes	Yes	Yes	Yes	Yes
Embedded configuration tools	Yes	Yes	Yes	Yes	Yes
Auto-Discovery	No	Yes	Yes	Yes	Yes
Remote OS deployment	No	Yes	Yes	Yes	Yes
Embedded driver pack	Yes	Yes	Yes	Yes	Yes
Full configuration inventory	Yes	Yes	Yes	Yes	Yes
Inventory export	Yes	Yes	Yes	Yes	Yes
Remote configuration	Yes	Yes	Yes	Yes	Yes
Zero-touch configuration	No	No	No	Yes	Yes
System Retire/Repurpose	Yes	Yes	Yes	Yes	Yes
Server Configuration Profile in GUI	Yes	Yes	Yes	Yes	Yes
Add BIOS configuration to iDRAC GUI	Yes	Yes	Yes	Yes	Yes
GPU properties	No	No	No	Yes	Yes
Diagnostics, Service, and	d Logging]			
Embedded diagnostic tools	Yes	Yes	Yes	Yes	Yes
Part Replacement	No	Yes	Yes	Yes	Yes
NOTE: After performin configuration, Lifecycle	ig part rep Logs rep	lacement of orts double	n RAID hardware, and t part replacement entri	the process is complete for replac es which is expected behavior.	ing firmware and
Easy Restore (system configuration)	Yes	Yes	Yes	Yes	Yes
Easy Restore Auto Timeout	Yes	Yes	Yes	Yes	Yes
i NOTE: Server Backup	and Resto	ore features	are not available in iDR	AC9 for PowerEdge Rx5xx/Cx5x	X.
LED Health status indicators	Yes	Yes	N/A	Yes	Yes
LCD screen (iDRAC9 requires optional)	Yes	Yes	N/A	Yes	Yes
iDRAC Quick Sync 2 (BLE/Wi-Fi hardware)	Yes	Yes	Yes	Yes	Yes



Feature	eature iDRAC iDRAC9 iDRAC9 Express iDRAC9 Enterprise iDRA 9 Basic Express for Blades iDRAC9 Enterprise Data		iDRAC9 Datacenter			
iDRAC Direct (front USB management port)	Yes	Yes	Yes	Yes Yes		
iDRAC Service Module (iSM) embedded	C Service Module Yes Yes Yes Yes Yes Yes		Yes			
iSM to in-band alert forwarding to consoles	Yes	Yes	Yes	Yes	Yes	
SupportAssist Collection (embedded)	Yes	Yes	Yes	Yes	Yes	
Crash screen capture	No	Yes	Yes	Yes	Yes	
Crash video capture ¹	No	No	No	Yes	Yes	
Agent Free Crash Video Capture (Windows only)	No	No	No	No	Yes	
Boot capture	No	No	No	Yes	Yes	
Manual reset for iDRAC (LCD ID button)	Yes	Yes	Yes	Yes	Yes	
Remote reset for iDRAC (requires iSM)	Yes	Yes	Yes	Yes	Yes	
Virtual NMI	Yes	Yes	Yes	Yes	Yes	
OS watchdog	Yes	Yes	Yes	Yes	Yes	
System Event Log	Yes	Yes	Yes	Yes	Yes	
Lifecycle Log	Yes	Yes	Yes	Yes	Yes	
Enhanced Logging in Lifecycle Controller Log	Yes	Yes	Yes	Yes	Yes	
Work notes	Yes	Yes	Yes	Yes	Yes	
Remote Syslog	No	No	No	Yes	Yes	
License management	Yes	Yes	Yes	Yes	Yes	
Improved Customer Experience						
iDRAC -Faster processor, more memory	N/A	Yes	N/A	Yes	Yes	
GUI rendered in HTML5	N/A	Yes	N/A	Yes	Yes	
Add BIOS configuration to iDRAC GUI	N/A	Yes	N/A	Yes	Yes	

[1] Requires iSM or OMSA agent on target server.

Interfaces and protocols to access iDRAC

The following table lists the interfaces to access iDRAC.



(i) NOTE: Using more than one interface at the same time may generate unexpected results.

Table 3. Interfaces and protocols to access iDRAC

Interface or Protocol	Description						
iDRAC Settings Utility (F2)	Use the iDRAC Settings utility to perform pre-OS operations. It has a subset of the features that are available in iDRAC web interface along with other features.						
	To access iDRAC Settings utility, press <f2> during boot and then click iDRAC Settings on the System Setup Main Menu page.</f2>						
Lifecycle Controller (F10)	Use Lifecycle Controller to perform iDRAC configurations. To access Lifecycle Controller, press <f10> during boot and go to System Setup > Advanced Hardware Configuration > iDRAC Settings. For more information, see <i>Lifecycle Controller User's Guide</i> available at dell.com/ idracmanuals.</f10>						
iDRAC Web Interface	Use the iDRAC web interface to manage iDRAC and monitor the managed system. The browser connects to the web server through the HTTPS port. Data streams are encrypted using 128-bit SSL to provide privacy and integrity. Any connection to the HTTP port is redirected to HTTPS. Administrators can upload their own SSL certificate through an SSL CSR generation process to secure the web server. The default HTTP and HTTPS ports can be changed. The user access is based on user privileges.						
OpenManage Enterprise (OME)	(i) NOTE: This interface is only available for MX platforms.						
Modular Web Interface	 In addition to monitoring and managing the chassis, use the OME-Modular web interface to: View the status of a managed system Update iDRAC firmware Configure iDRAC network settings Log in to iDRAC web interface 						
	Start, stop, or reset the managed system						
	Update BIOS, PERC, and supported network adapters						
	available at https://www.dell.com/openmanagemanuals.						
CMC Web Interface	(i) NOTE: This interface is not available in MX platforms.						
	 In addition to monitoring and managing the chassis, use the CMC web interface to: View the status of a managed system Update iDRAC firmware Configure iDRAC network settings Log in to iDRAC web interface Start, stop, or reset the managed system Update BIOS, PERC, and supported network adapters 						
Server LCD Panel/ Chassis LCD Panel	 Use the LCD on the server front panel to: View alerts, iDRAC IP or MAC address, user programmable strings. Set DHCP Configure iDRAC static IP settings. 						
	For blade servers, the LCD is on the chassis front panel and is shared between all the blades.						
	To reset iDRAC without rebooting the server, press and hold the System Identification button $m heta$ for 16 seconds.						
	() NOTE: LCD panel is only available with rack or tower systems that support front bezel. For blade servers, the LCD is on the chassis front panel and is shared between all the blades.						
RACADM	 Use this command-line utility to perform iDRAC and server management. You can use RACADM locally and remotely. Local RACADM command-line interface runs on the managed systems that have Server Administrator installed. Local RACADM communicates with iDRAC through its in-band IPMI host 						

Inserido ao protocolo **22.951.206-4** por: **Pedro Henrique Golin Linhares** em: 31/03/2025 12:49. A autenticidade deste documento pode ser validada no endereço: https://www.eprotocolo.pr.gov.br/spiweb/validarDocumento com o código: fc1c44dfe0686d5547ffd691199be699.



Table 3. Interfaces and protocols to access iDRAC (continued)

Interface or Protocol	Description					
	 interface. Since it is installed on the local managed system, users are required to log in to the operating system to run this utility. A user must have a full administrator privilege or be a root user to use this utility. Remote RACADM is a client utility that runs on a management station. It uses the out-of-band network interface to run RACADM commands on the managed system and uses the HTTPs channel. The -r option runs the RACADM command over a network. Firmware RACADM is accessible by logging in to iDRAC using SSH. You can run the firmware RACADM commands without specifying the iDRAC IP, user name, or password. You do not have to specify the iDRAC IP, user name, or password to run the firmware RACADM commands. After you enter the RACADM prompt, you can directly run the commands without the racadm prefix. 					
iDRAC RESTful API and Redfish	The Redfish Scalable Platforms Management API is a standard defined by the Distributed Management Task Force (DMTF). Redfish is a next-generation systems management interface standard, which enables scalable, secure, and open server management. It is a new interface that uses RESTful interface semantics to access data that is defined in model format to perform out- of-band systems management. It is suitable for a wide range of servers ranging from stand-alone servers to rack mount and bladed environments and for large scale cloud environments. Redfish provides the following benefits over existing server management methods: Increased simplicity and usability High data security Programmable interface that can be easily scripted Follows widely-used standards					
	For iDRAC Redfish API guide, go to www.api-marketplace.com					
WSMan	The LC-Remote Service is based on the WSMan protocol to do one-to-many systems management tasks. You must use WSMan client such as WinRM client (Windows) or the OpenWSMan client (Linux) to use the LC-Remote Services functionality. You can also use Power Shell or Python to script to the WSMan interface.					
	protocol used for systems management. iDRAC uses WSMan to convey Distributed Management Task Force (DMTF) Common Information Model (CIM)–based management information. The CIM information defines the semantics and information types that can be modified in a managed system. The data available through WSMan is provided by iDRAC instrumentation interface mapped to the DMTF profiles and extension profiles.					
	 For more information, see the following: Lifecycle Controller Remote Services Quick Start Guide available at https://www.dell.com/ idracmanuals. MOFs and Profiles — http://downloads.dell.com/wsman. DMTF website — dmtf.org/standards/profiles 					
SSH	Use SSH to run RACADM commands. The SSH service is enabled by default on iDRAC. The SSH service can be disabled in iDRAC. iDRAC only supports SSH version 2 with the RSA host key algorithm. A unique 1024-bit RSA host key is generated when you power-up iDRAC for the first time.					
IPMITool	Use the IPMITool to access the remote system's basic management features through iDRAC. The interface includes local IPMI, IPMI over LAN, IPMI over Serial, and Serial over LAN. For more information on IPMITool, see the <i>Dell OpenManage Baseboard Management Controller Utilities User's Guide</i> at dell.com/idracmanuals.					
NTLM	iDRAC allows NTLM to provide authentication, integrity, and confidentiality to the users. NT LAN Manager (NTLM) is a suite of Microsoft security protocols and it works in a Windows network.					
SMB	iDRAC9 supports the Server Message Block (SMB) Protocol. This is a network file sharing protocol and the default minimum SMB version supported is 2.0, SMBv1 is no longer supported.					



Table 3. Interfaces and protocols to access iDRAC (continued)

Interface or Protocol	Description
NFS	iDRAC9 supports Network File System (NFS). This is a distributed filesystem protocol that enables users to mount remote directories on the servers.

iDRAC port information

The following table lists the ports that are required to remotely access iDRAC through firewall. These are the default ports iDRAC listens to for connections. Optionally, you can modify most of the ports. To modify ports, see Configuring services.

Table 4. Ports iDRAC listens for connections

Port number	Туре	Function	Configurabl e port	Maximum Encryption Level				
22	TCP	SSH	Yes	256-bit SSL				
80	TCP	HTTP	Yes	None				
161	UDP	SNMP Agent	Yes	None				
443	TCP	 Web GUI access with HTTPS Virtual Console and Virtual Media with eHTML5 option Virtual Console and Virtual Media with HTML5 option when web server redirection is enabled 	Yes	256-bit SSL				
623	UDP	RMCP/RMCP+	No	128-bit SSL				
5000	TCP	iDRAC to iSM	No	256-bit SSL				
(i) NOTE: Maximum encryption level is 256-bit SSL if both iSM 3.4 or higher and iDRAC firmware 3.30.30.30 or higher are installed.								
5900	TCP	Virtual console and virtual media with HTML5, Java and ActiveX option	Yes	128-bit SSL				
5901	TCP	VNC	Yes	128-bit SSL				
	i NOTE: Port 5901 opens when VNC feature is enabled.							

The following table lists the ports that iDRAC uses as a client:

Table 5. Ports iDRAC uses as client

Port number	Туре	Function	Configurable port	Maximum Encryption Level
25	TCP	SMTP	Yes	None
53	UDP	DNS	No	None
68	UDP	DHCP-assigned IP address	No	None
69	TFTP	TFTP	No	None
123	UDP	Network Time Protocol (NTP)	No	None
162	UDP	SNMP trap	Yes	None
445	TCP	Common Internet File System (CIFS)	No	None
636	TCP	LDAP Over SSL (LDAPS)	No	256-bit SSL
2049	TCP	Network File System (NFS)	No	None



Table 5. Ports iDRAC uses as client (continued)

Port number	Туре	Function	Configurable port	Maximum Encryption Level	
3269 TCP LDAPS for global catalog (GC)		No	256-bit SSL		
5353	UDP	mDNS	No	None	
() NOTE: When node initiated discovery or Group Manager is enabled, iDRAC uses mDNS to communicate through port 5353. However, when both are disabled, port 5353 is blocked by iDRAC's internal firewall and appears as open filtered port in the port scans.					
514	UDP	Remote syslog	Yes	None	

Other documents you may need

Some of the iDRAC interfaces have the integrated *Online Help* document that can be accessed by clicking on the help (?) icon. The *Online Help* provides detailed information about the fields available on the web interface and the descriptions for the same. In addition, following documents are available on the Dell Support website at **dell.com/support** that provide additional information about the setup and operation of iDRAC in your system.

- The iDRAC Redfish API Guide available at https://developer.dell.com provides information about Redfish API.
- The *iDRAC RACADM CLI Guide* provides information about the RACADM sub-commands, supported interfaces, and iDRAC property database groups and object definitions.
- The Systems Management Overview Guide provides brief information about the various software available to perform systems management tasks.
- The *Dell Remote Access Configuration Tool User's Guide* provides information on how to use the tool to discover iDRAC IP addresses in your network and perform one-to-many firmware updates and active directory configurations for the discovered IP addresses.
- The *Dell Systems Software Support Matrix* provides information about the various Dell systems, the operating systems supported by these systems, and the Dell OpenManage components that can be installed on these systems.
- The iDRAC Service Module User's Guide provides information to install the iDRAC Service Module.
- The Dell OpenManage Server Administrator Installation Guide contains instructions to help you install Dell OpenManage Server Administrator.
- The Dell OpenManage Management Station Software Installation Guide contains instructions to help you install Dell OpenManage management station software that includes Baseboard Management Utility, DRAC Tools, and Active Directory Snap-In.
- The Dell OpenManage Baseboard Management Controller Management Utilities User's Guide has information about the IPMI interface.
- The *Release Notes* provides last-minute updates to the system or documentation or advanced technical reference material intended for experienced users or technicians.

The following system documents are available to provide more information:

- The safety instructions that came with your system provide important safety and regulatory information. For additional regulatory information, see the Regulatory Compliance home page at **dell.com/regulatory_compliance**. Warranty information may be included within this document or as a separate document.
- The Rack Installation Instructions included with your rack solution describe how to install your system into a rack.
- The *Getting Started Guide* provides an overview of system features, setting up your system, and technical specifications.
- The *Installation and Service Manual* provides information about system features and describes how to troubleshoot the system and install or replace system components.

Contacting Dell

NOTE: If you do not have an active Internet connection, you can find contact information on your purchase invoice, packing slip, bill, or Dell product catalog.

Dell provides several online and telephone-based support and service options. Availability varies by country and product, and some services may not be available in your area. To contact Dell for sales, technical support, or customer service issues, visit https://www.dell.com/contactdell.



Accessing documents from Dell support site

You can access the required documents in one of the following ways:

- Using the following links:
 - For all Enterprise Systems Management and OpenManage Connections documents https://www.dell.com/ esmmanuals
 - For OpenManage documents https://www.dell.com/openmanagemanuals
 - For iDRAC and Lifecycle Controller documents https://www.dell.com/idracmanuals
 - For Serviceability Tools documents https://www.dell.com/serviceabilitytools
 - For Client Command Suite Systems Management documents https://www.dell.com/omconnectionsclient

Accessing documents using the product search

- 1. Go to https://www.dell.com/support.
- 2. In the Enter a Service Tag, Serial Number... search box, type the product name. For example, PowerEdge or iDRAC.

A list of matching products is displayed.

- 3. Select your product and click the search icon or press enter.
- 4. Click **DOCUMENTATION**.
- 5. Click MANUALS AND DOCUMENTS.

Accessing documents using product selector

You can also access documents by selecting your product.

- 1. Go to https://www.dell.com/support.
- 2. Click Browse all products.
- 3. Click the desired product category, such as Servers, Software, Storage, and so on.
- 4. Click the desired product and then click the desired version if applicable.

(i) NOTE: For some products, you may need to navigate through the subcategories.

- 5. Click **DOCUMENTATION**.
- 6. Click MANUALS AND DOCUMENTS.

Accessing Redfish API Guide

The Redfish API guide is now available at the Dell API Marketplace. To access the Redfish API guide:

- 1. Go to www.api-marketplace.com.
- 2. Click Explore API and then click APIs.
- 3. Under iDRAC9 Redfish API, click View More.



Logging in to iDRAC

You can log in to iDRAC as an iDRAC user, a Microsoft Active Directory user, or a Lightweight Directory Access Protocol (LDAP) user. You can also log in using OpenID Connect and Single Sign-On or Smart Card.

To improve security, each system is shipped with a unique password for iDRAC, which is available on the system information tag. This unique password improves security of iDRAC and your server. The default user name is *root*.

While ordering the system, you can choose to retain the legacy password—calvin—as the default password. If you choose to retain the legacy password, the password is not available on the system information tag.

In this version, DHCP is enabled by default and iDRAC IP address is assigned dynamically.

() NOTE:

- You must have login to iDRAC privilege to log in to iDRAC.
- iDRAC GUI does not support browser buttons such as **Back**, **Forward**, or **Refresh**.

NOTE: For information about recommended characters for user names and passwords, see Recommended characters in user names and passwords.

To change the default password, see Changing the default login password.

Customizable security banner

You can customize the security notice that is displayed on the login page. You can use SSH, RACADM, Redfish, or WSMan to customize the notice. Depending on the language you use, the notice can be either 1024 or 512 UTF-8 characters long.

OpenID Connect

(i) NOTE: This feature is only available for MX platforms.

You can log in to iDRAC using credentials of other web consoles such as Dell EMC OpenManage Enterprise (OME) - Modular. When this feature is enabled, the console starts managing the user permissions on the iDRAC. iDRAC provides the user session with all the permissions that are specified by the console.

(i) NOTE: When lockdown mode is enabled, OpenID Connect login options are not displayed in iDRAC login page.

You can now get access to detailed help without logging in to iDRAC. Use the links on the iDRAC login page to access help and version information, drivers and downloads, manuals and TechCenter.

Topics:

- Force Change of Password (FCP)
- Logging into iDRAC using OpenID Connect
- Logging in to iDRAC as local user, Active Directory user, or LDAP user
- Logging in to iDRAC as a local user using a smart card
- Logging in to iDRAC using Single Sign-On
- Accessing iDRAC using remote RACADM
- Accessing iDRAC using local RACADM
- Accessing iDRAC using firmware RACADM
- Simple 2-Factor Authentication (Simple 2FA)
- RSA SecurID 2FA
- Viewing system health
- Logging in to iDRAC using public key authentication
- Multiple iDRAC sessions

Logging in to iDRAC 37



- Secure default password
- Changing the default login password
- Enabling or disabling default password warning message
- Password Strength Policy
- IP Blocking
- Enabling or disabling OS to iDRAC Pass-through using web interface
- Enabling or disabling alerts using RACADM

Force Change of Password (FCP)

The 'Force Change of Password' feature prompts you to change the factory default password of the device. The feature can be enabled as part of factory configuration.

The FCP screen appears after successful user authentication and cannot be skipped. Only after the user enters a password, normal access and operation will be allowed. The state of this attribute will not be affected by a 'Reset Configuration to Defaults' operation.

(i) NOTE: To set or reset the FCP attribute, you must have Login privilege and User configuration privilege.

(i) NOTE: When FCP is enabled, 'Default Password Warning' setting is disabled after changing the default user password.

(i) NOTE: When root user logs in via Public Key Authentication (PKA), FCP is bypassed.

When FCP is enabled, following actions are not allowed:

- Login to iDRAC through any UI except IPMIpover-LAN interface which uses CLI with default user credentials.
- Login to iDRAC through OMM app via Quick Sync-2
- Add a member iDRAC in Group Manager.

Logging into iDRAC using OpenID Connect

(i) NOTE: This feature is only available in MX platforms.

To log in to iDRAC using the OpenID Connect:

- In a supported web browser, type https://[iDRAC-IP-address] and press Enter. The Login page is displayed.
- 2. Select OME Modular from the Log In with: menu. The console login page is displayed.
- 3. Enter the console User name and Password.
- 4. Click Log in.

You are logged in to iDRAC with the console user privileges.

(i) NOTE: When lockdown mode is enabled, OpenID Connect login option is not displayed in iDRAC login page.

Logging in to iDRAC as local user, Active Directory user, or LDAP user

Before you log in to iDRAC using the web interface, ensure that you have configured a supported web browser and the user account is created with the required privileges.

(i) NOTE: The user name is not case-sensitive for an Active Directory user. The password is case-sensitive for all users.

(i) NOTE: In addition to Active Directory, openLDAP, openDS, Novell eDir, and Fedora-based directory services are supported.

(i) NOTE: LDAP authentication with OpenDS is supported. The DH key must be larger than 768 bits.



NOTE: RSA feature can be configured and enabled for LDAP user, but the RSA does not support if the LDAP is configured on Microsoft active directory. Hence LDAP user login fails. RSA is supported only for OpenLDAP.

To log in to iDRAC as local user, Active Directory user, or LDAP user:

- 1. Open a supported web browser.
- 2. In the Address field, type https://[iDRAC-IP-address] and press Enter.
 - () NOTE: If the default HTTPS port number (port 443) changes, enter: https://[iDRAC-IP-address]:[port-number] where [iDRAC-IP-address] is the iDRAC IPv4 or IPv6 address and [port-number] is the HTTPS port number.

The **Login** page is displayed.

- 3. For a local user:
 - In the **Username** and **Password** fields, enter your iDRAC user name and password.
 - From the **Domain** drop-down menu, select **This iDRAC**.
- 4. For an Active Directory user, in the User name and Password fields, enter the Active Directory user name and password. If you have specified the domain name as a part of the username, select This iDRAC from the drop-down menu. The format of the user name can be: <domain>\<username>, <domain>\<username>, or <user>@<domain>.

For example, dell.com\john_doe, or JOHN_DOE@DELL.COM.

If the domain is not specified in the user name, select the Active Directory domain from the **Domain** drop-down menu.

- 5. For an LDAP user, in the **Username** and **Password** fields, enter your LDAP user name and password. Domain name is not required for LDAP login. By default, **This iDRAC** is selected in the drop-down menu.
- 6. Click **Submit**. You are logged in to iDRAC with the required user privileges.

If you log in with Configure Users privileges and the default account credentials, and if the default password warning feature is enabled, the **Default Password Warning** page is displayed allowing you to easily change the password.

Logging in to iDRAC as a local user using a smart card

Before you log in as a local user using Smart Card, make sure to:

- Upload user smart card certificate and the trusted Certificate Authority (CA) certificate to iDRAC.
- Enable smart card logon.

The iDRAC web interface displays the smart card logon page for users who are configured to use the smart card.

- (i) **NOTE:** Depending on the browser settings, you are prompted to download and install the smart card reader ActiveX plug-in when using this feature for the first time.
- To log in to iDRAC as a local user using a smart card:
- 1. Access the iDRAC web interface using the link https://[IP address].

The **iDRAC Login** page is displayed prompting you to insert the smart card.

- **NOTE:** If the default HTTPS port number (port 443) changes, type: https://[IP address]:[port number] where, [IP address] is the IP address for the iDRAC and [port number] is the HTTPS port number.
- 2. Insert the smart card into the reader and click **Login**.
 - A prompt is displayed for the smart card's PIN. A password is not required.
- 3. Enter the Smart Card PIN for local smart card users.
 - You are logged in to the iDRAC.
 - (i) **NOTE:** If you are a local user for whom **Enable CRL check for Smart Card Logon** is enabled, iDRAC attempts to download the certificate revocation list (CRL) and checks the CRL for the user's certificate. The login fails if the certificate is listed as revoked in the CRL or if the CRL cannot be downloaded for some reason.

(i) NOTE: If you log in to iDRAC using smart card when RSA is enabled, RSA token is bypassed and you can login directly.



Logging in to iDRAC as an Active Directory user using a smart card

Before you log in as an Active Directory user using smart card, ensure that you:

- Upload a Trusted Certificate Authority (CA) certificate (CA-signed Active Directory certificate) to iDRAC.
- Configure the DNS server.
- Enable Active Directory login.
- Enable smart card login.

To log in to iDRAC as an Active Directory user using smart card:

1. Log in to iDRAC using the link https://[IP address].

The **iDRAC Login** page is displayed prompting you to insert the smart card.

NOTE: If the default HTTPS port number (port 443) is changed, type: https://[IP address]:[port number] where, [IP address] is the iDRAC IP address and [port number] is the HTTPS port number.

- 2. Insert the smart card and click Login.
 - A prompt is displayed for the smart card's **PIN**.
- 3. Enter the PIN and click **Submit**.

You are logged in to iDRAC with your Active Directory credentials.

(i) NOTE:

If the smart card user is present in Active Directory, an Active Directory password is not required.

Logging in to iDRAC using Single Sign-On

When Single Sign-On (SSO) is enabled, you can log in to iDRAC without entering your domain user authentication credentials, such as user name and password.

(i) NOTE: When AD user configures SSO while RSA is enabled, the RSA token is bypassed and user logs in directly.

Logging in to iDRAC SSO using iDRAC web interface

Before logging in to iDRAC using Single Sign-On, ensure that:

- You have logged in to your system using a valid Active Directory user account.
- Single Sign-On option is enabled during Active Directory configuration.
- To log in to iDRAC using web interface:
- 1. Log in to your management station using a valid Active Directory account.
- 2. In a web browser, type https://[FQDN address].
 - () NOTE: If the default HTTPS port number (port 443) has been changed, type: https://[FQDN address]:[port number] where [FQDN address] is the iDRAC FQDN (iDRACdnsname.domain. name) and [port number] is the HTTPS port number.
 - (i) NOTE: If you use IP address instead of FQDN, SSO fails.

iDRAC logs you in with appropriate Microsoft Active Directory privileges, using your credentials that were cached in the operating system when you logged in using a valid Active Directory account.

Logging in to iDRAC SSO using CMC web interface

(i) NOTE: This feature is not available on MX platforms.

Using the SSO feature, you can launch the iDRAC web interface from the CMC web interface. A CMC user has the CMC user privileges when launching iDRAC from CMC. If the user account is present in CMC and not in iDRAC, the user can still launch iDRAC from CMC.

If iDRAC network LAN is disabled (LAN Enabled = No), SSO is not available.

40 Logging in to iDRAC



If the server is removed from the chassis, iDRAC IP address is changed, or there is a problem in iDRAC network connection, the option to Launch iDRAC is grayed-out in the CMC web interface.

For more information, see the Chassis Management Controller User's Guide available at https://www.dell.com/cmcmanuals.

Accessing iDRAC using remote RACADM

You can use remote RACADM to access iDRAC using RACADM utility.

For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

If the management station has not stored the iDRAC's SSL certificate in its default certificate storage, a warning message is displayed when you run the RACADM command. However, the command is executed successfully.

() NOTE: The iDRAC certificate is the certificate iDRAC sends to the RACADM client to establish the secure session. This certificate is either issued by a CA or self-signed. In either case, if the management station does not recognize the CA or signing authority, a warning is displayed.

Validating CA certificate to use remote RACADM on Linux

Before running remote RACADM commands, validate the CA certificate that is used for secure communications.

To validate the certificate for using remote RACADM:

1. Convert the certificate in DER format to PEM format (using openssl command-line tool):

openssl x509 -inform pem -in [yourdownloadedderformatcert.crt] -outform pem -out [outcertfileinpemformat.pem] -text

- Find the location of the default CA certificate bundle on the management station. For example, for RHEL5 64 bit, it is /etc/pki/tls/cert.pem.
- 3. Append the PEM formatted CA certificate to the management station CA certificate. For example, use the cat command: cat testcacert.pem >> cert.pem
- **4.** Generate and upload the server certificate to iDRAC.

Accessing iDRAC using local RACADM

For information to access iDRAC using local RACADM, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/ idracmanuals.

Accessing iDRAC using firmware RACADM

You can use SSH interface to access iDRAC and run firmware RACADM commands. For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Simple 2-Factor Authentication (Simple 2FA)

iDRAC offers simple 2-factor authentication option to enhance the security to the local users for logging in. When you log in from a source IP-address which is different from the last login, you will be prompted to enter the second factor authentication details.

Simple two factor authentication has two steps of authentication:

- iDRAC User name and password
- Simple 6 digit code which is sent to the user via email. User needs to enter this 6 digit code when prompted at login.

() NOTE:

• To receive 6 digit code, it is mandatory to configure 'Custom Sender Address' and have valid SMTP configuration.



- The 2FA code expires after 10 minutes or is invalidated if it is already consumed before expiry.
- If a user attempts to login from another location with a different IP-Address while a pending 2FA challenge for the original IP-Address is still outstanding, the same token will be sent for login attempt from the new IP address.
- The feature is supported with iDRAC Enterprise or Datacenter license.

When 2FA is enabled, following actions are not allowed:

- Login to iDRAC through any UI which uses CLI with default user credentials.
- Login to iDRAC through OMM app via Quick Sync-2
- Add a member iDRAC in Group Manager.

NOTE: Racadm, Redfish, WSMAN, IPMI LAN, Serial, CLI from a source IP works only after successful login from supported interfaces like iDRAC GUI and SSH.

RSA SecurID 2FA

iDRAC can be configured to authenticate with a single RSA AM server at a time. The global settings on RSA AM server apply to all iDRAC local users, AD, and LDAP users.

(i) NOTE: RSA SecurID 2FA teature is available only on Datacenter license.

Following are the pre-requisites before you configure iDRAC to enable RSA SecurID:

- Configure Microsoft Active Directory server.
- If you try to enable RSA SecurID on all AD users, add the AD server to the RSA AM server as an Identity Source.
- Ensure you have a generic LDAP server.
- For all LDAP users, the Identity Source to the LDAP server must be added in RSA AM server.

To enable RSA SecurID on iDRAC, the following attributes from the RSA AM server are required:

- 1. RSA Authentication API URL The URL syntax is: https://<rsa-am-server-hostname>:<port>/mfa/v1_1, and by default the port is 5555.
- RSA Client-ID By default, the RSA client ID is the same as the RSA AM server hostname. Find the RSA client ID at RSA AM server's authentication agent configuration page.
- 3. RSA Access Key The Access Key can be retreived on RSA AM by navigating to Setup
 > System Settings > RSA SecurID > Authetication APIsection, which is usually displayed as

198cv5x195fdi86u43jw0q069byt0x37umlfwxc2gnp4s0xk11ve21ffum4s8302. To configure the settings through iDRAC GUI:

- Go to iDRAC Settings > Users.
- From Local Users section, select an existing local user and click Edit.
- Scroll down to the bottom of the Configuration page.
- In RSA SecurID section, Click the link RSA SecurID Configuration to view or edit these settings.

You can also configure the settings as follows:

- Go to **iDRAC Settings** > Users.
- From Directory Services section, select Microsoft Active Service or Generic LDAP Directory Service, and click Edit.
- In RSA SecurID section, Click the link RSA SecurID Configuration to view or edit these settings.

4. RSA AM server certificate (chain)

You can login to iDRAC using RSA SecurID token via iDRAC GUI and SSH.

RSA SecurID Token App

You need to install RSA SecurID Token app on you system or on smart phone. When you try to log in to iDRAC, you are asked to input the passcode shown in the app.

If a wrong passcode is entered, the RSA AM server challenges the user to provide the "Next Token." This may happen even though the user may have entered the correct passcode. This entry proves that the user owns the right Token that generates the right passcode.



You get the **Next Token** from RSA SecurID Token app by clicking on **Options**. Check **Next Token**, and the next passcode is available. Time is critical in this step. Otherwise, iDRAC may fail the verification of the next token. If the iDRAC user login session times out, it requires another attempt to log in

If a wrong passcode is entered, the RSA AM server will challenge the user to provide the "Next Token." This challenge happens even though the user may have later entered the correct passcode. This entry proves that the user owns the right Token that generates the right passcodes.

To get the next token from RSA SecurID Token app, click on **Options** and check **Next Token**. A new token is generated. Time is critical in this step. Otherwise, iDRAC may fail the verification of the next token. If the iDRAC user login session times out, it requires another attempt to log in.

Viewing system health

Before you perform a task or trigger an event, you can use RACADM to check if the system is in a suitable state. To view the remote service status from RACADM, use the getremoteservicesstatus command.

Table 6. Possible values for system status

	Host System	Lifecycle Controller (LC)		Real Time Status		Overall Status
• • • •	Powered Off In POST Out of POST Collecting System Inventory Automated Task Execution Lifecycle Controller Unified Server Configurator Server has halted at F1/F2 error prompt because of a POST error Server has halted at F1/F2/F11 prompt because there are no bootable devices available Server has entered F2 setup menu Server has entered F11 Boot Manager menu	 Ready Not Initialized Reloading data Disabled In Recovery In Use 	•	Ready Not Ready	•	Ready Not Ready
1. 2. 3. 4.	1. Read/Write: Read Only 2. User Privilege: Login User 3. License Required: iDRAC Express or iDRAC Enterprise 4. Dependency: None					

Logging in to iDRAC using public key authentication

You can log in to the iDRAC over SSH without entering a password. You can also send a single RACADM command as a command line argument to the SSH application. The command line options behave like remote RACADM since the session ends after the command is completed.

For example:

Logging in:

ssh username@<domain>

or

ssh username@<IP_address>

Fis. <u>1086</u> Mov. <u>33</u>

where IP address is the IP address of the iDRAC.

Sending RACADM commands:

ssh username@<domain> racadm getversion

```
ssh username@<domain> racadm getsel
```

Multiple iDRAC sessions

The following table provides the number of iDRAC sessions that are possible using the various interfaces.

Table 7. Multiple iDRAC sessions

Interface	Number of Sessions
iDRAC Web Interface	8
Remote RACADM	4
Firmware RACADM	SSH - 4
	Serial - 1

iDRAC allows multiple sessions for the same user. After a user has created the maximum number of allowed sessions, other users cannot log in to the iDRAC. This can cause a *Denial of Service* for a legitimate administrator user.

In case of session exhaustion, perform the following remedies:

- If webserver-based sessions are exhausted, you can still login via SSH or local RACADM.
- An administrator can then terminate existing sessions using racadm commands (racadm getssninfo; racadm closessn -i <index>).

Secure default password

All supported systems are shipped with a unique default password for iDRAC, unless you choose to set *calvin* as the password while ordering the system. The unique password helps improve the security of iDRAC and your server. To further enhance security, it is recommended that you change the default password.

The unique password for your system is available on the system information tag. To locate the tag, see the documentation for your server at https://www.dell.com/support.

(i) NOTE: For PowerEdge C6420, M640, and FC640, the default password is calvin.

NOTE: Resetting iDRAC to the factory default settings reverts the default password to the one that the server was shipped with.

If you have forgotten the password and do not have access to the system information tag, there are a few methods to reset the password locally or remotely.

Resetting default iDRAC password locally

If you have physical access to the system, you can reset the password using the following:

- iDRAC Setting utility (System Setup)
- Local RACADM
- OpenManage Mobile
- Server management USB port
- USB-NIC



Resetting default password using the iDRAC Settings utility

You can access the iDRAC settings utility using the System Setup of your server. Using the iDRAC reset to defaults all feature, you can reset the iDRAC login credentials to default.

WARNING: Resetting iDRAC to default all, resets the iDRAC to the factory defaults.

To reset iDRAC using iDRAC Settings utility:

- 1. Reboot the server and press <F2>.
- 2. In the System Setup page, click iDRAC Settings.
- 3. Click Reset iDRAC configurations to defaults all.
- 4. Click **Yes** to confirm, and then click **Back**.
- 5. Click Finish.

The server restarts after all iDRAC settings are set to default settings.

Resetting default password using local RACADM

- 1. Log in to the host OS installed on the system.
- 2. Access the local RACADM interface.
- 3. Follow the instructions in Changing the default login password using RACADM.

Resetting default password using OpenManage Mobile

You can use the OpenManage Mobile (OMM) to log in and change the default password. To log in to iDRAC using OMM, scan the QR code on the system information tag. For more information about using OMM, see the OMM documentation at *OME* - *Modular for PowerEdge MX7000 Chassis User's Guide* available at https://www.dell.com/openmanagemanuals.

them from the default values, enter the updated credentials.

Resetting default password using the server management USB port

(i) **NOTE:** These steps require that the USB management port is enabled and configured.

Using Server Configuration Profile file

Create a Server Configuration Profile (SCP) file with a new password for the default account, place it on a memory key, and use the server management USB port on the server to upload the SCP file. For more information on creating the file, see Using USB port for server management.

Accessing iDRAC using a laptop

Connect a laptop to the server management USB port and access iDRAC to change the password. For more information, see Accessing iDRAC interface over direct USB connection.

Changing default password using USB-NIC

If you have access to a keyboard, mouse, and a display device, connect to the server using the USB-NIC to access the iDRAC interface and change the default password.

- 1. Connect the devices to the system.
- 2. Use a supported browser to access the iDRAC interface using the iDRAC IP.
- **3.** Follow the instructions in Changing the default login password using web interface.



Resetting default iDRAC password remotely

If you do not have physical access to the system, you can reset the default password remotely.

Remote — Provisioned system

If you have an operating system installed on the system, use a remote desktop client to log in to the server. After you log into the server, use any of the local interfaces such as RACADM or web interface to change the password.

Remote — Non-provisioned system

If there is no operating system installed on the server and if you have a PXE setup available, use PXE and then use RACADM to reset the password.

Changing the default login password

The warning message that allows you to change the default password is displayed if:

- You log in to iDRAC with Configure User privilege.
- The default password warning feature is enabled.
- The default iDRAC user name and password are provided on the system information tag.

A warning message is also displayed when you log in to iDRAC using SSH, remote RACADM, or the Web interface. For Web interface and SSH, a single warning message is displayed for each session. For remote RACADM, the warning message is displayed for each command.

NOTE: For information on recommended characters for user names and passwords, see Recommended characters in user names and passwords.

Changing the default login password using web interface

When you log in to the iDRAC web interface, if the **Default Password Warning** page is displayed, you can change the password. To do this:

- 1. Select the Change Default Password option.
- 2. In the New Password field, enter the new password.

NOTE: For information on recommended characters for user names and passwords, see Recommended characters in user names and passwords.

- 3. In the **Confirm Password** field, enter the password again.
- 4. Click Continue.

The new password is configured and you are logged in to iDRAC.

(i) NOTE: Continue is enabled only if the passwords entered in the New Password and Confirm Password fields match.

For information about the other fields, see the *iDRAC Online Help*.

Changing the default login password using RACADM

To change the password, run the following RACADM command:

racadm set iDRAC.Users.<index>.Password <Password>

where, <index> is a value from 1 to 16 (indicates the user account) and <password> is the new user defined password. (i) NOTE: The index for the default account is 2.

For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

46 Logging in to iDRAC



NOTE: For information on recommended characters for user names and passwords, see Recommended characters in user names and passwords.

Changing the default login password using iDRAC settings utility

To change the default login password using iDRAC settings utility:

- 1. In the iDRAC Settings utility, go to User Configuration.
- The **iDRAC Settings User Configuration** page is displayed. **2.** In the **Change Password** field, enter the new password.

(i) NOTE: For information on recommended characters for user names and passwords, see Recommended characters in user names and passwords.

3. Click **Back**, click **Finish**, and then click **Yes**. The details are saved.

Enabling or disabling default password warning message

You can enable or disable the display of the default password warning message. To do this, you must have Configure Users privilege.

Password Strength Policy

Using iDRAC interface, you can check the password strength policy and check any errors if the policy is not met. The password policy cannot be applied to previously saved passwords, Server Configuration Profiles (SCP) copied from other servers, and embedded passwords in the profile.

To access Password settings, go to iDRAC Settings > Users > Password Settings.

Following fields are available in this section:

- Minimum Score Specifies the minimum password strength policy score. The values in this field are:
 - \circ 0 No protection
 - 1 Weak protection
 - 2 Medium protection
 - 3 Strong protection
- Simple Policy Specifies the required characters in a secure password. It has the following options:
 - Upper Case Letters
 - Numbers
 - Symbols
 - Minimum Length
- **Regular Expression** The Regular expression along with the Minimum score is used for password enforcement. The values are 1-4.

IP Blocking

You can use IP blocking to dynamically determine when excessive login failures occur from an IP address and block or prevent the IP address from logging into the iDRAC9 for a preselected time span. IP blocking includes:

- The number of allowable login failures.
- The timeframe in seconds when these failures must occur.
- The amount of time, in seconds, when the IP address is prevented from establishing a session after the total allowable number of failures is exceeded.



As consecutive login failures accumulate from a specific IP address, they are tracked by an internal counter. When the user logs in successfully, the failure history is cleared and the internal counter is reset.

NOTE: When consecutive login attempts are refused from the client IP address, some SSH clients may display the following message:

ssh exchange identification: Connection closed by remote host

(i) NOTE: IP blocking feature supports upto 5 IP ranges. You can see / set these only via RACADM.

Table 8. Login Retry Restriction Properties

Property	Definition		
iDRAC.IPBlocking.BlockEnable	Enables the IP blocking feature. When consecutive failures		
	iDRAC.IPBlocking.FailCount		
	from a single IP address are encountered within a specific amount of time		
	iDRAC.IPBlocking.FailWindow		
	all further attempts to establish a session from that address are rejected for a certain timespan		
	iDRAC.IPBlocking.PenaltyTime		
iDRAC.IPBlocking.FailCount	Sets the number of login failures from an IP address before the login attempts are rejected.		
iDRAC.IPBlocking.FailWindow	The time, in seconds during which the failed attempts are counted. When the failures occur beyond this time period, the counter gets reset.		
iDRAC.IPBlocking.PenaltyTime	Defines the timespan in seconds when all login attempts from an IP address with excessive failures are rejected.		

Enabling or disabling OS to iDRAC Pass-through using web interface

To enable OS to iDRAC Pass-through using Web interface:

- 1. Go to iDRAC Settings > Connectivity > Network > OS to iDRAC Pass-through. The OS to iDRAC Pass-through page is displayed.
- 2. Change the State to **Enabled**.
- **3.** Select any of the following options for Pass-through Mode:
 - LOM The OS to iDRAC pass-through link between the iDRAC and the host operating system is established through the LOM or NDC.
 - **USB NIC** The OS to iDRAC pass-through link between the iDRAC and the host operating system is established through the internal USB bus.

(i) NOTE: If you set the pass-through mode to LOM, ensure that:

- OS and iDRAC are on the same subnet
- NIC selection in Network Settings is set to LOM
- 4. If the server is connected in shared LOM mode, then the OS IP Address field is disabled.



NOTE: If the VLAN is enabled on the iDRAC, the LOM-Passthrough will only function in shared LOM mode with VLAN tagging configured on the host.

() NOTE:

- When Pass-through mode is set to LOM, it is not possible to launch iDRAC from host OS after cold boot.
- We have purposefully removed the LOM Pass-through using Dedicated mode feature.
- 5. If you select **USB NIC** as the pass-through configuration, enter the IP address of the USB NIC.

The default value is 169.254.1.1. It is recommended to use the default IP address. However, if this IP address conflicts with an IP address of other interfaces of the host system or the local network, you must change it.

Do not enter 169.254.0.3 and 169.254.0.4 IPs. These IPs are reserved for the USB NIC port on the front panel when an A/A cable is used.

- **NOTE:** If IPv6 is preferred, the default address is fde1:53ba:e9a0:de11::1. If needed, this address can be modified in the idrac.OS-BMC.UsbNicULA setting. If IPv6 is not wanted on the USB-NIC, it can be disabled by changing the address to "::"
- 6. Click Apply.
- 7. Click **Test Network Configuration** to check if the IP is accessible and the link is established between the iDRAC and the host operating system.

Enabling or disabling alerts using RACADM

Use the following command:

```
racadm set iDRAC.IPMILan.AlertEnable <n>
```

n=0 — Disabled

n=1 — Enabled



Setting up managed system

If you need to run local RACADM or enable Last Crash Screen capture, install the following from the *Dell Systems Management Tools and Documentation* DVD:

- Local RACADM
- Server Administrator

For more information about Server Administrator, see *OpenManage Server Administrator User's Guide* available at https://www.dell.com/openmanagemanuals.

Topics:

- Setting up iDRAC IP address
- Modifying local administrator account settings
- Setting up managed system location
- Optimizing system performance and power consumption
- Setting up management station
- Configuring supported web browsers
- Updating device firmware
- Viewing and managing staged updates
- Rolling back device firmware
- Monitoring iDRAC using other Systems Management tools
- Support Server Configuration Profile Import and Export
- Secure Boot Configuration from BIOS Settings or F2
- BIOS recovery

Setting up iDRAC IP address

You must configure the initial network settings based on your network infrastructure to enable the communication to and from iDRAC. You can set up the IP address using one of the following interfaces:

- iDRAC Settings utility
- Lifecycle Controller (see Lifecycle Controller User's Guide)
- Chassis or Server LCD panel (see the system's Installation and Service Manual)
 - **NOTE:** On blade servers, you can configure the network settings using the chassis LCD panel only during the initial CMC configuration. You cannot reconfigure iDRAC using the Chassis LCD panel after the chassis is deployed.
 - CMC Web interface (not applicable for MX platforms) (see Chassis Management Controller User's Guide)

In case of rack and tower servers, you can set up the IP address or use the default iDRAC IP address 192.168.0.120 to configure initial network settings, including setting up DHCP or the static IP for iDRAC.

In case of blade servers, the iDRAC network interface is disabled by default.

After you configure iDRAC IP address:

- Ensure that you change the default user name and password.
- Access iDRAC through any of the following interfaces:
 - iDRAC Web interface using a supported browser (Internet Explorer, Firefox, Chrome, or Safari)
 - Secure Shell (SSH) Requires a client such as PuTTY on Windows. SSH is available by default in most of the Linux systems and hence does not require a client.
 - IPMITool (uses IPMI command) or shell prompt (requires Dell customized installer in Windows or Linux, available from *Systems Management Documentation and Tools* DVD or https://www.dell.com/support)



Setting up iDRAC IP using iDRAC settings utility

To set up the iDRAC IP address:

- 1. Turn on the managed system.
- 2. Press <F2> during Power-on Self-test (POST).
- 3. In the System Setup Main Menu page, click iDRAC Settings. The iDRAC Settings page is displayed.
- 4. Click Network. The Network page is displayed.
- **5.** Specify the following settings:
 - Network Settings
 - Common Settings
 - IPv4 Settings
 - IPv6 Settings
 - IPMI Settings
 - VLAN Settings
- 6. Click **Back**, click **Finish**, and then click **Yes**. The network information is saved and the system reboots.

Configuring the network settings

To configure the network settings:

(i) NOTE: For information about the options, see the *iDRAC* Settings Utility Online Help.

- 1. Under Enable NIC, select Enabled.
- 2. From the NIC Selection drop-down menu, select one of the following ports based on the network requirement:

(i) NOTE: This option is not available in MX platforms.

• **Dedicated** — Enables the remote access device to use the dedicated network interface available on the Remote Access Controller (RAC). This interface is not shared with the host operating system and routes the management traffic to a separate physical network, enabling it to be separated from the application traffic.

This option implies that iDRAC's dedicated network port routes its traffic separately from the server's LOM or NIC ports. The Dedicated option allows iDRAC to be assigned an IP address from the same subnet or different subnet in comparison to the IP addresses assigned to the Host LOM or NICs to manage the network traffic.

(i) **NOTE:** In blade servers, the Dedicated option is displayed as **Chassis (Dedicated)**.

- LOM1
- LOM2
- LOM3
- LOM4
- () NOTE: In the case of rack and tower servers, two LOM options (LOM1 and LOM2) or all four LOM options are available depending on the server model. In blade servers with two NDC ports, two LOM options (LOM1 and LOM2) are available and on servers with four NDC ports, all four LOM options are available.
- **NOTE:** Shared LOM is not supported on the *Intel 2P X520–k bNDC 10 G* if they are used in a full-height server with two NDCs because they do not support hardware arbitration.
- **3.** From the **NIC Selection** drop-down menu, select the port from which you want to access the system remotely, following are the options:
 - (i) NOTE: This feature is not available in MX platforms.
 - **NOTE:** You can select either the dedicated network interface card or from a list of LOMs available in the Quad port or Dual port mezzanine cards.



• **Chassis (Dedicated)**: Enables the remote access device to use the dedicated network interface available on the Remote Access Controller (RAC). This interface is not shared with the host operating system and routes the management traffic to a separate physical network, enabling it to be separated from the application traffic.

This option implies that iDRAC's dedicated network port routes its traffic separately from the server's LOM or NIC ports. The Dedicated option allows iDRAC to be assigned an IP address from the same subnet or different subnet in comparison to the IP addresses assigned to the Host LOM or NICs to manage the network traffic.

• For Quad port cards—LOM1-LOM16

• For Dual port cards—LOM1, LOM2, LOM5, LOM6, LOM9, LOM10, LOM13, LOM14.

4. From the **Failover Network** drop-down menu, select one of the remaining LOMs. If a network fails, the traffic is routed through the failover network.

For example, to route the iDRAC network traffic through LOM2 when LOM1 is down, select **LOM1** for **NIC Selection** and **LOM2** for **Failover Network**.

INOTE: This option is disabled if **NIC Selection**is set to **Dedicated**.

NOTE: When using the **Failover network** settings, it is recommended that all the LOM ports to be connected to the same network.

For more details, refer to the section Modifying network settings using web interface

- Under Auto Negotiation, select On if iDRAC must automatically set the duplex mode and network speed.
 This option is available only for dedicated mode. If enabled, iDRAC sets the network speed to 10, 100, or 1000 Mbps based on the network speed.
- 6. Under Network Speed, select either 10 Mbps or 100 Mbps.
 - **NOTE:** You cannot manually set the Network Speed to 1000 Mbps. This option is available only if **Auto Negotiation** option is enabled.
- 7. Under Duplex Mode, select Half Duplex or Full Duplex option.

(i) NOTE: This option is disabled if Auto Negotiation is set to Enabled.

- () NOTE: If network teaming is configured for the host OS using the same network adapter as NIC Selection, then the Failover Network should also be configured. NIC Selection and Failover Network should use the ports that are configured as a part of the network team. If more than two ports are used as part of the network team, then the Failover Network selection should be "All".
- 8. Under NIC MTU, enter the Maximum Transmission Unit size on the NIC.
 - () NOTE: The default and maximum limit for MTU on NIC is 1500, and the minimum value is 576. An MTU value of 1280 or greater is required if IPv6 is enabled.

Common settings

If network infrastructure has DNS server, register iDRAC on the DNS. These are the initial settings requirements for advanced features such as Directory services—Active Directory or LDAP, Single Sign On, and smart card.

To register iDRAC:

- 1. Enable Register DRAC on DNS.
- 2. Enter the DNS DRAC Name.
- 3. Select Auto Config Domain Name to automatically acquire domain name from DHCP. Else, provide the DNS Domain Name.

For **DNS iDRAC Name** field, the default name format is *idrac-Service_Tag*, where Service_Tag is the service tag of the server. The maximum length is 63 characters and the following characters are supported:

- A-Z
- a-z
- 0-9
- Hyphen (-)



Configuring the IPv4 settings

To configure the IPv4 settings:

1. Select Enabled option under Enable IPv4.

(i) **NOTE:** In the 14th generation of the PowerEdge servers, DHCP is enabled by default.

- 2. Select **Enabled** option under **Enable DHCP**, so that DHCP can automatically assign the IP address, gateway, and subnet mask to iDRAC. Else, select **Disabled** and enter the values for:
 - Static IP Address
 - Static Gateway
 - Static Subnet Mask
- Optionally, enable Use DHCP to obtain DNS server address, so that the DHCP server can assign the Static Preferred DNS Server and Static Alternate DNS Server. Else, enter the IP addresses for Static Preferred DNS Server and Static Alternate DNS Server.

Configuring the IPv6 settings

Based on the infrastructure setup, you can use IPv6 address protocol.

To configure the IPv6 settings:

- **NOTE:** If IPv6 is set to static, ensure that you configure the IPv6 gateway manually, which is not needed in case of dynamic IPV6. Failing to configure manually in case of static IPV6 results in loss of communication.
- 1. Select Enabled option under Enable IPv6.
- 2. For the DHCPv6 server to automatically assign the IP address, gateway, and subnet mask to iDRAC, select **Enabled** option under **Enable Auto-configuration**.

(i) NOTE: You can configure both static IP and DHCP IP at the same time.

- 3. In the Static IP Address 1 box, enter the static IPv6 address.
- 4. In the Static Prefix Length box, enter a value between 0 and 128.
- 5. In the Static Gateway box, enter the gateway address.
 - (i) NOTE: If you configure static IP, the current IP address 1 displays static IP and the IP address 2 displays dynamic IP. If you clear the static IP settings, the current IP address 1 displays dynamic IP.
- 6. If you are using DHCP, enable **DHCPv6 to obtain DNS Server addresses** to obtain Primary and Secondary DNS server addresses from DHCPv6 server. You can configure the following if required:
 - In the Static Preferred DNS Server box, enter the static DNS server IPv6 address.
 - In the Static Alternate DNS Server box, enter the static alternate DNS server.

Configuring the IPMI settings

To enable the IPMI Settings:

- 1. Under Enable IPMI Over LAN, select Enabled.
- 2. Under Channel Privilege Limit, select Administrator, Operator, or User.
- **3.** In the **Encryption Key** box, enter the encryption key in the format 0 to 40 hexadecimal characters (without any blanks characters.) The default value is all zeros.

VLAN settings

You can configure iDRAC into the VLAN infrastructure. To configure VLAN settings, perform the following steps:

- (i) **NOTE:** On blade servers that are set as **Chassis (Dedicated)**, the VLAN settings are read-only and can be changed only using CMC. If the server is set in shared mode, you can configure VLAN settings in shared mode in iDRAC.
- 1. Under Enable VLAN ID, select Enabled.



- 2. In the VLAN ID box, enter a valid number from 1 to 4094.
- 3. In the **Priority** box, enter a number from 0 to 7 to set the priority of the VLAN ID.

(i) NOTE: After enabling VLAN, the iDRAC IP is not accessible for some time.

Setting up iDRAC IP using the CMC web interface

To set up the iDRAC IP address using the Chassis Management Controller (CMC) Web interface:

NOTE: You must have Chassis Configuration Administrator privilege to set up iDRAC network settings from CMC. The CMC option is applicable only for blade servers.

- 1. Log in to the CMC Web interface.
- 2. Go to iDRAC Settings > Settings > CMC. The Deploy iDRAC page is displayed.
- 3. Under **iDRAC Network Settings**, select **Enable LAN** and other network parameters as per requirements. For more information, see *CMC online help*.
- For additional network settings specific to each blade server, go to Server Overview > <server name>. The Server Status page is displayed.
- 5. Click Launch iDRAC and go to iDRAC Settings > Connectivity > Network.
- 6. In the **Network** page, specify the following settings:
 - Network Settings
 - Common Settings
 - IPV4 Settings
 - IPV6 Settings
 - IPMI Settings
 - VLAN Settings
 - Advanced Network Settings

(i) NOTE: For more information, see *iDRAC Online Help*.

7. To save the network information, click Apply.

For more information, see the Chassis Management Controller User's Guide available at https://www.dell.com/cmcmanuals.

Auto-discovery

The Auto-discovery feature allows newly installed servers to automatically discover the remote management console that hosts the provisioning server. The provisioning server provides custom administrative user credentials to iDRAC so that the unprovisioned server can be discovered and managed from the management console. For more information about provisioning server, see the *Lifecycle Controller Remote Services Quick Start Guide* available at https://www.dell.com/idracmanuals.

Provisioning server works with a static IP address. Auto-discovery feature on the iDRAC is used to find the provisioning server using DHCP/Unicast DNS/mDNS.

- When iDRAC has the console address, it sends its own service tag, IP address, Redfish port number, Web certificate etc.
- This information is periodically published to consoles.

DHCP, DNS server, or the default DNS host name discovers the provisioning server. If DNS is specified, the provisioning server IP is retrieved from DNS and the DHCP settings are not required. If the provisioning server is specified, discovery is skipped so neither DHCP nor DNS is required.

Auto-discovery can be enabled using the following ways:

1. Using iDRAC GUI: iDRAC Settings > Connectivity > iDRAC Auto Discover



2. Using RACADM:

Transfer (F, St. Sender), A. St. Sender St. St. St. Sender (F, St. Sender), Sender (F, St. Sender), St. Sender (F, St. Sender), St. Sender (F, St. Sender), St. Sender St. St. Sender (F, St. Sender), St. St. Sender St. St. Sender (F, St. Sender), St. S	Approximation .	
land af Lington Construction Co		a de com a de composition de la composition de la composition de la composition de la composition entre de la composition de la composition entre de la composition de la composition entre de la composition de la compositione
na de l'Angenetante martin de la construcción de l	Contraction of the second seco	- Martin and Constraints (1997) and a second strategy (2011) - N Strategy (1998) (2011) - Martin Theorem (1997)
Dell'Al Conference and Conference	nations and the second	mark,
nani a selektrongennen der i dag of arreitet ervelt at en der er ta selekt. Nervel i stand – er en den er en en er en en er en er er en er	and the Polyment of the second	 Market Market Constraints Market Constraints Market Constraints
	naal a natio hargenn Nage at it name Darretter a	andar - Lan d'Arméni (1998) a la altern la satta 1995 - Lan Sona Constanti (1998) a constanti (1998) a constanti (1998) a constanti (1998) 1999 - Manuala 1999 - Manuala 1999 - Manuala



To enable provisioning server using iDRAC Settings utility:

- **1.** Turn on the managed system.
- During POST, press F2, and go to iDRAC Settings > Remote Enablement. The iDRAC Settings Remote Enablement page is displayed.
- 3. Enable auto-discovery, enter the provisioning server IP address, and click **Back**.

NOTE: Specifying the provisioning server IP is optional. If it is not set, it is discovered using DHCP or DNS settings (step 7).

4. Click Network.

The iDRAC Settings Network page is displayed.

- 5. Enable NIC.
- 6. Enable IPv4.

(i) NOTE: IPv6 is not supported for auto-discovery.

7. Enable DHCP and get the domain name, DNS server address, and DNS domain name from DHCP.

(i) NOTE: Step 7 is optional if the provisioning server IP address (step 3) is provided.

Configuring servers and server components using Auto Config

The Auto Config feature configures and provisions all the components in a server in a single operation. These components include BIOS, iDRAC, and PERC. Auto Config automatically imports a Server Configuration Profile (SCP) XML or JSON file containing all configurable parameters. The DHCP server that assigns the IP address also provides the details for accessing the SCP file.

SCP files are created by configuring a gold configuration server. This configuration is then exported to a shared NFS, CIFS, HTTP or HTTPS network location that is accessible by the DHCP server and the iDRAC of the server being configured. The SCP file name can be based on the Service Tag or model number of the target server or can be given as a generic name. The DHCP server uses a DHCP server option to specify the SCP file name (optionally), SCP file location, and the user credentials to access the file location.

When the iDRAC obtains an IP address from the DHCP server that is configured for Auto Config, iDRAC uses the SCP to configure the server's devices. Auto Config is invoked only after the iDRAC gets its IP address from the DHCP server. If it does not get a response or an IP address from the DHCP server, then Auto Config is not invoked.

HTTP and HTTPS file sharing options are supported for iDRAC firmware 3.00.00.00 or later. Details of the HTTP or HTTPS address need to be provided. In case the proxy is enabled on the server, the user needs to provide further proxy settings to allow HTTP or HTTPS to transfer information. The –s option flag is updated as:

Table 9. Different Share Types and pass in values

-s (ShareType)	pass in
NFS	0 or nfs
CIFS	2 or cifs
HTTP	5 or http
HTTPS	6 or https

(i) NOTE: HTTPS certificates are not supported with Auto Config. Auto Config ignores certificate warnings.

Following list describes the required and optional parameters to pass in for the string value:

-f (Filename): name of exported Server Configuration Profile file. This is required for iDRAC firmware versions prior to 2.20.20.20.

-n (Sharename): name of network share. This is required for NFS or CIFS.

-s (ShareType): pass in either 0 for NFS, 2 for CIFS, 5 for HTTP and 6 for HTTPS. This is a mandatory field for iDRAC firmware versions 3.00.00.00.

-i (IPAddress): IP address of the network share. This is a mandatory field.

56 Setting up managed system


- -u (Username) : username that has access to network share. This is a mandatory field for CIFS.
- -p (Password): user password that has access to network share. This is a mandatory field for CIFS.
- -d (ShutdownType): either 0 for graceful or 1 for forced (default setting: 0). This is an optional field.
- -t (Timetowait): time to wait for the host to shutdown (default setting: 300). This is an optional field.

-e (EndHostPowerState): either 0 for OFF or 1 for ON (default setting 1). This is an optional field.

The additional option flags are supported in iDRAC firmware 3.00.00.00 or later to enable the configuration of HTTP proxy parameters and set the retry timeout for accessing the Profile file:

- -pd (ProxyDefault): Use default proxy setting. This is an optional field.
- -pt (ProxyType): The user can pass in http or socks (default setting http). This is an optional field.
- -ph (ProxyHost): IP address of the proxy host. This is an optional field.
- -pu (ProxyUserName): username that has access to the proxy server. This is required for proxy support.
- -pp (ProxyPassword): user password that has access to the proxy server. This is required for proxy support.
- -po (ProxyPort): port for the proxy server (default setting is 80). This is an optional field.
- -to (Timeout): specifies the retry timeout in minutes for obtaining config file (default is 60 minutes).

For iDRAC firmware 3.00.00.00 or later, JSON format Profile files are supported. The following file names will be used if the Filename parameter is not present:

- <service tag>-config.xml, Example: CDVH7R1-config.xml
- <model number>-config.xml, Example: R640-config.xml
- config.xml
- <service tag>-config.json, Example:CDVH7R1-config.json
- <model number>-config.json, Example: R630-config.json
- config.json

NOTE: More information about HTTP can be found in the 14G Support for HTTP and HTTPS across IDRAC9 with Lifecycle Controller Interfaces white paper at https://www.dell.com/support.

(i) NOTE:

- Auto Config can only be enabled when DHCPv4 and the Enable IPV4 options are enabled.
- Auto Config and Auto Discovery features are mutually exclusive. Disable Auto Discovery for Auto Config to work.
- The Auto Config is disabled after a server has carried out an Auto Config operation.

If all the Dell PowerEdge servers in the DHCP server pool are of the same model type and number, then a single SCP file (config.xml) is required. The config.xml file name is used as the default SCP file name. In addition to .xml file, .json files can also be used with 14G systems. The file can be config.json.

The user can configure individual servers requiring different configuration files mapped using individual server Service Tags or server models. In an environment that has different servers with specific requirements, different SCP file names can be used to distinguish each server or server type. For example, if there are two server models to configure — PowerEdge R740s and PowerEdge R540s, use two SCP files, R740-config.xml and R540-config.xml.

NOTE: iDRAC server configuration agent automatically generates the configuration filename using the server Service Tag, model number, or the default filename — config.xml.

NOTE: If none of these files are on the network share, then the server configuration profile import job is marked as failed for file not found.

Auto Config sequence

- 1. Create or modify the SCP file that configures the attributes of Dell servers.
- 2. Place the SCP file in a share location that is accessible by the DHCP server and all the Dell servers that are assigned IP address from the DHCP server.
- 3. Specify the SCP file location in vendor-option 43 field of DHCP server.
- 4. The iDRAC while acquiring IP address advertises vendor class identifier. (Option 60)



- 5. The DHCP server matches the vendor class to the vendor option in the dhcpd.conf file and sends the SCP file location and, if specified the SCP file name to the iDRAC.
- 6. The iDRAC processes the SCP file and configures all the attributes listed in the file.

DHCP options

DHCPv4 allows many globally defined parameters to be passed to the DHCP clients. Each parameter is known as a DHCP option. Each option is identified with an option tag, which is a 1-byte value. Option tags 0 and 255 are reserved for padding and end of options, respectively. All other values are available for defining options.

The DHCP Option 43 is used to send information from the DHCP server to the DHCP client. The option is defined as a text string. This text string is set to contain the values of the SCP filename, share location and the credentials to access the location. For example,

```
option myname code 43 = text;
subnet 192.168.0.0 netmask 255.255.255.0 {
    # default gateway
        option routers 192.168.0.1;
    option subnet-mask 255.255.255.0;
    option nis-domain "domain.org";
    option domain-name "domain.org";
    option domain-name-servers 192.168.1.1;
    option domain-name-servers 192.168.1.1;
    option time-offset -18000; #Eastern Standard Time
    option vendor-class-identifier "iDRAC";
    set vendor-string = option vendor-class-identifier;
    option myname "-f system_config.xml -i 192.168.0.130 -u user -p password -n cifs -s
2 -d 0 -t 500";
```

where, -i is the location of the Remote File Share and -f is the file name in the string along with the credentials to the Remote File Share.

The DHCP Option 60 identifies and associates a DHCP client with a particular vendor. Any DHCP server configured to take action based on a client's vendor ID should have Option 60 and Option 43 configured. With Dell PowerEdge servers, the iDRAC identifies itself with vendor ID: *iDRAC*. Therefore, you must add a new 'Vendor Class' and create a 'scope option' under it for 'code 60,' and then enable the new scope option for the DHCP server.

Configuring option 43 on Windows

To configure option 43 on Windows:

- 1. On the DHCP server, go to **Start** > **Administration Tools** > **DHCP** to open the DHCP server administration tool.
- 2. Find the server and expand all items under it.
- **3.** Right-click on **Scope Options** and select **Configure Options**. The **Scope Options** dialog box is displayed.
- 4. Scroll down and select 043 Vendor Specific Info.
- 5. In the **Data Entry** field, click anywhere in the area under **ASCII** and enter the IP address of the server that has the share location, which contains the SCP file.

The value appears as you type it under the ASCII, but it also appears in binary to the left.

6. Click **OK** to save the configuration.

Configuring option 60 on Windows

To configure option 60 on Windows:

- 1. On the DHCP server, go to Start > Administration Tools > DHCP to open the DHCP server administration tool.
- 2. Find the server and expand the items under it.
- 3. Right-click on IPv4 and choose Define Vendor Classes.
- 4. Click Add.
 - A dialog box with the following fields is displayed:
 - Display name:
 - Description:
 - ID: Binary: ASCII:



- 5. In the **Display name:** field, type iDRAC.
- 6. In the **Description:** field, type Vendor Class.
- 7. Click in the **ASCII:** section and type iDRAC.
- 8. Click OK and then Close.
- 9. On the DHCP window, right-click IPv4 and select Set Predefined Options.
- 10. From the Option class drop-down menu, select iDRAC (created in step 4) and click Add.
- 11. In the **Option Type** dialog box, enter the following information:
 - Name iDRAC
 - Data Type String
 - Code 060
 - Description Dell vendor class identifier
- 12. Click OK to return to the DHCP window.
- 13. Expand all items under the server name, right-click Scope Options and select Configure Options.
- 14. Click the Advanced tab.
- 15. From the Vendor class drop-down menu, select **iDRAC**. The 060 iDRAC is displayed in the Available Options column.
- 16. Select 060 iDRAC option.
- 17. Enter the string value that must be sent to the iDRAC (along with a standard DHCP provided IP address). The string value helps in importing the correct SCP file.

For the option's **DATA entry, String Value** setting, use a text parameter that has the following letter options and values:

- Filename (-f) Indicates the name of the exported Server Configuration Profile(SCP) file.
- Sharename (-n) Indicates the name of the network share.
- ShareType (-s) —

Alongside supporting NFS and CIFS-based file sharing, iDRAC firmware 3.00.00.00 or later also supports accessing profile files by using HTTP and HTTPS. The -s option flag is updated as follows:

-s (ShareType): type nfs or 0 for NFS; cifs or 2 for CIFS; http or 5 for HTTP; or https or 6 for HTTPS (mandatory).

- IPAddress (-i) Indicates the IP address of the file share.
 NOTE: Sharename (-n), ShareType (-s), and IPAddress (-i) are required attributes that must be passed. -n is not required for HTTP or HTTPs.
- Username (-u) Indicates the user name required to access the network share. This information is required only for CIFS.
- Password (-p) Indicates the password required to access the network share. This information is required only for CIFS.
- ShutdownType (-d) Indicates the mode of shutdown. 0 indicates Graceful shutdown and 1 indicates Forced shutdown.

i NOTE: The default setting is 0.

- Timetowait (-t) Indicates the time the host system waits before shutting down. The default setting is 300.
- EndHostPowerState (-e) Indicates the power state of the host. 0 indicates OFF and 1 indicates ON. The default setting is 1.

(i) NOTE: ShutdownType (-d), Timetowait (-t), and EndHostPowerState (-e) are optional attributes.

NFS: -f system_config.xml -i 192.168.1.101 -n /nfs_share -s 0 -d 1 **CIFS**: -f system_config.xml -i 192.168.1.101 -n cifs_share -s 2 -u *<USERNAME>* -p *<PASSWORD>* -d 1 -t 400 **HTTP**: -f system_config.json -i 192.168.1.101 -s 5 **HTTP**: -f http_share/system_config.xml -i 192.168.1.101 -s http **HTTP**: -f system_config.xml -i 192.168.1.101 -s http -n http_share **HTTPS**: -f system_config.json -i 192.168.1.101 -s http

Configuring option 43 and option 60 on Linux

Update the /etc/dhcpd.conf file. The steps to configure the options are similar to the steps for Windows:

1. Set aside a block or pool of addresses that this DHCP server can allocate.



2. Set the option 43 and use the name vendor class identifier for option 60.

```
option myname code 43 = text;
subnet 192.168.0.0 netmask 255.255.0.0 {
#default gateway
    option routers
                                   192.168.0.1;
                                    255.255.255.0;
    option subnet-mask
    option domain-name
option domai
                                  "domain.org";
                                  "domain.org";
    option domain-name-servers
                                        192.168.1.1;
    option time-offset
                                    -18000;
                                                # Eastern Standard Time
    option vendor-class-identifier "iDRAC";
  set vendor-string = option vendor-class-identifier;
option myname "-f system_config.xml -i 192.168.0.130 -u user -p password -n cifs -s 2 -d 0 -t 500";
    range dynamic-bootp 192.168.0.128 192.168.0.254;
    default-lease-time 21600;
    max-lease-time 43200;
         }
}
```

The following are the required and optional parameters that must be passed in the vendor class identifier string:

• Filename (-f) — Indicates the name of the exported Server Configuration Profile file.

(i) NOTE: For more information on file naming rules, see Configuring servers and server components using Auto Config.

- Sharename (-n) Indicates the name of the network share.
- ShareType (-s) Indicates the share type. 0 indicates NFS, 2 indicates CIFS, 5 indicates HTTP, and 6 indicates HTTPS.
 NOTE: Example for Linux NFS, CIFS, HTTP, HTTPS share:
 - NFS:-f system_config.xml -i 192.168.0.130 -n /nfs -s 0 -d 0 -t 500

Ensure that you use NFS2 or NFS3 for NFS network share.

- CIFS: -f system_config.xml -i 192.168.0.130 -n sambashare/config_files -s 2 -u user -p password -d 1 -t 400
- HTTP:-f system_config.xml -i 192.168.1.101 -s http -n http_share
- HTTPS: -f system config.json -i 192.168.1.101 -s https
- IPAddress (-i) Indicates the IP address of the file share.

NOTE: Sharename (-n), ShareType (-s), and IPAddress (-i) are required attributes that must be passed. -n is not required for HTTP or HTTPS.

- Username (-u) Indicates the user name required to access the network share. This information is required only for CIFS.
- Password (-p) Indicates the password required to access the network share. This information is required only for CIFS.
- ShutdownType (-d) Indicates the mode of shutdown. 0 indicates Graceful shutdown and 1 indicates Forced shutdown.
 (i) NOTE: The default setting is 0.
- Timetowait (-t) Indicates the time the host system waits before shutting down. The default setting is 300.
- EndHostPowerState (-e) Indicates the power state of the host. 0 indicates OFF and 1 indicates ON. The default
 setting is 1.

(i) NOTE: ShutdownType (-d), Timetowait (-t), and EndHostPowerState (-e) are optional attributes.

The following is an example of a static DHCP reservation from a dhcpd.conf file:

```
host my_host {
host my_host {
hardware ethernet b8:2a:72:fb:e6:56;
fixed-address 192.168.0.211;
option host-name "my_host";
option myname " -f r630_raid.xml -i 192.168.0.1 -n /nfs -s 0 -d 0 -t 300";
}
```

(i) NOTE: After editing the dhcpd.conf file, make sure to restart the dhcpd service to apply the changes.



Prerequisites before enabling Auto Config

Before enabling the Auto config feature, make sure that following are already set:

- Supported network share (NFS, CIFS, HTTP and HTTPS) is available on the same subnet as the iDRAC and DHCP server. Test the network share to ensure that it can be accessed and that the firewall and user permissions are set correctly.
- Server configuration profile is exported to the network share. Also, make sure that the necessary changes in the SCP file are complete so that proper settings can be applied when the Auto Config process is initiated.
- DHCP server is set and the DHCP configuration is updated as required for iDRAC to call the server and initiate the Auto Config feature.

Enabling Auto Config using iDRAC web interface

Make sure that DHCPv4 and the Enable IPv4 options are enabled and Auto-discovery is disabled.

To enable Auto Config:

- 1. In the iDRAC web interface, go to **iDRAC Settings** > **Connectivity** > **Network** > **Auto Config**. The **Network** page is displayed.
- 2. In the Auto Config section, select one of the following options from the Enable DHCP Provisioning drop-down menu:
 - Enable Once Configures the component only once using the SCP file referenced by the DHCP server. After this, Auto Config is disabled.
 - **Enable once after reset** After the iDRAC is reset, configures the components only once using the SCP file referenced by the DHCP server. After this, Auto Config is disabled.
 - **Disable** Disables the Auto Config feature.
- **3.** Click **Apply** to apply the setting. The network page automatically refreshes.

Enabling Auto Config using RACADM

To enable Auto Config feature using RACADM, use the iDRAC.NIC.AutoConfig object.

For more information, see the iDRAC RACADM CLI Guide available at https://www.dell.com/idracmanuals.

For more information on the Auto Config feature, see the Zero-Touch, bare-metal server provisioning using the Dell EMC iDRAC with Lifecycle Controller Auto Config feature white paper available at the https://www.dell.com/support.

Using hash passwords for improved security

On PowerEdge servers with iDRAC version 3.00.00.00, you can set user passwords and BIOS passwords using a one-way hash format. The user authentication mechanism is not affected (except for SNMPv3 and IPMI) and you can provide the password in plain text format.

With the new password hash feature:

• You can generate your own SHA256 hashes to set iDRAC user passwords and BIOS passwords. This allows you to have the SHA256 values in the server configuration profile, RACADM, and WSMan. When you provide the SHA256 password values, you cannot authenticate through SNMPv3 and IPMI.

NOTE: Remote RACADM or WSMan or Redfish cannot be used for Hash password Configuration / Replacement for IDRAC. You can use SCP for Hash Password Configuration / Replacement on Remote RACADM or WSMan or Redfish.

- You can set up a template server including all the iDRAC user accounts and BIOS passwords using the current plain text
 mechanism. After the server is set up, you can export the server configuration profile with the password hash values. The
 export includes the hash values required for SNMPv3 and IPMI authentication. After importing this profile, you must use the
 the latest Dell IPMI tool, if you use an older tool, the IPMI authentication fails for the users who have the hashed password
 values set.
- The other interfaces such as IDRAC GUI will show the user accounts enabled.

You can generate the hash password with and without Salt using SHA256.

You must have Server Control privileges to include and export hash passwords.

If access to all accounts is lost, use iDRAC Settings Utility or local RACADM and perform reset iDRAC to default task.

If the password of the iDRAC user account is set with the SHA256 password hash only and not the other hashes (SHA1v3Key or MD5v3Key or IPMIKey), then authentication through SNMP v3 and IPMI is not available.

Fls.<u>1104</u> Mov. 33

Hash password using RACADM

To set hash passwords, use the following objects with the set command:

- iDRAC.Users.SHA256Password
- iDRAC.Users.SHA256PasswordSalt

() NOTE: SHA256Password and SHA256PasswordSalt fields are reserved for XML import and do not set them using command line tools. Setting one of the fields can potentially lock out the current user from logging into iDRAC. When a password is imported using SHA256Password, the iDRAC will not be enforcing the password length check.

Use the following command to include the hash password in the exported server configuration profile:

racadm get -f <file name> -l <NFS / CIFS / HTTP / HTTPS share> -u <username> -p
<password> -t <filetype> --includePH

You must set the Salt attribute when the associated hash is set.

(i) NOTE: The attributes are not applicable to the INI configuration file.

Hash password in server configuration profile

The new hash passwords can be optionally exported in the server configuration profile.

When importing server configuration profile, you can uncomment the existing password attribute or the new password hash attribute(s). If both are uncommented an error is generated and the password is not set. A commented attribute is not applied during an import.

Generating hash password without SNMPv3 and IPMI authentication

Hash password can be generated without SNMPv3 and IPMI authentication with or without salt. Both require SHA256.

To generate hash password with salt:

1. For the iDRAC user accounts, you must salt the password using SHA256.

When you salt the password, a 16-bytes binary string is appended. The Salt is required to be 16 bytes long, if provided. Once appended, it becomes a 32 character string. The format is "password"+"salt", for example:

Password = SOMEPASSWORD

Salt = ALITTLEBITOFSALT-16 characters are appended

2. Open a Linux command prompt, and run the following command:

Generate Hash-> echo-n SOMEPASSWORDALITTLEBITOFSALT|sha256sum -><HASH> Generate Hex Representation of Salt -> echo -n ALITTLEBITOFSALT | xxd -p -> <HEX-

SALT>

set iDRAC.Users.4.SHA256Password <HASH>

set iDRAC.Users.4.SHA256PasswordSalt <HEX-SALT>

3. Provide hash value and salt in the imported server configuration profile, the RACADM commands, Redfish, or WSMan.



NOTE: If you wish to clear a previously salted password, then ensure that the password-salt is explicitly set to an empty string i.e.

```
set iDRAC.Users.4.SHA256Password
ca74e5fe75654735d3b8d04a7bdf5dcdd06f1c6c2a215171a24e5a9dcb28e7a2
```

set iDRAC.Users.4.SHA256PasswordSalt

 After setting the password, the normal plain text password authentication works except that SNMP v3 and IPMI authentication fails for the iDRAC user accounts that had passwords updated with hash.

Modifying local administrator account settings

After setting the iDRAC IP address, you can modify the local administrator account settings (that is, user 2) using the iDRAC Settings utility. To do this:

- In the iDRAC Settings utility, go to User Configuration. The iDRAC Settings User Configuration page is displayed.
- 2. Specify the details for User Name, LAN User Privilege, Serial Port User Privilege, and Change Password. For information about the options, see the *iDRAC Settings Utility Online Help*.
- **3.** Click **Back**, click **Finish**, and then click **Yes**. The local administrator account settings are configured.

Setting up managed system location

You can specify the location details of the managed system in the data center using the iDRAC Web interface or iDRAC Settings utility.

Setting up managed system location using web interface

To specify the system location details:

- In the iDRAC web interface, go to System > Details > System Details. The System Details page is displayed.
- **2.** Under **System Location**, enter the location details of the managed system in the data center. For information about the options, see the *iDRAC Online Help*.
- 3. Click **Apply**. The system location details are saved in iDRAC.

Setting up managed system location using RACADM

To specify the system location details, use the System.Location group objects.

For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Setting up managed system location using iDRAC settings utility

To specify the system location details:

- In the iDRAC Settings utility, go to System Location.
 The iDRAC Settings System Location page is displayed.
- 2. Enter the location details of the managed system in the data center. For information about the options, see the *iDRAC* Settings Utility Online Help.
- **3.** Click **Back**, click **Finish**, and then click **Yes**. The details are saved.



Optimizing system performance and power consumption

The power required to cool a server can contribute a significant amount to the overall system power. Thermal control is the active management of system cooling through fan speed and system power management to make sure that the system is reliable while minimizing system power consumption, airflow, and system acoustic output. You can adjust the thermal control settings and optimize against the system performance and performance-per-Watt requirements.

Using the iDRAC Web interface, RACADM, or the iDRAC Settings Utility, you can change the following thermal settings:

- Optimize for performance
- Optimize for minimum power
- Set the maximum air exhaust temperature
- Increase airflow through a fan offset, if required

• Increase airflow through increasing minimum fan speed

Following are the list of features in thermal management:

- System Airflow Consumption: Displays the real-time system airflow consumption (in CFM), allowing airflow balancing at rack and datacenter level.
- Custom Delta-T: Limit air temperature rise from inlet air to exhaust to right-size your infrastructure level cooling.
- Exhaust Temperature Control: Specify the temperature limit of the air exiting the server to match your datacenter needs.
- **Custom PCIe inlet temperature**: Choose the right input inlet temperature to match 3rd party device requirements.
- **PCIe Airflow settings**: Provides a comprehensive PCIe device cooling view of the server and allows cooling customization of 3rd party cards.

Modifying thermal settings using iDRAC web interface

To modify the thermal settings:

- 1. In the iDRAC Web interface, go to Configuration > System Settings > Hardware Settings > Cooling Configuration.
- 2. Specify the following:
 - Thermal Profile Optimization Select the thermal profile:
 - **Default Thermal Profile Settings (Minimum Power)** Implies that the thermal algorithm uses the same system profile settings that is defined under **System BIOS** > **System BIOS Settings** > **System Profile Settings** page.

By default, this option is set to **Default Thermal Profile Settings**. You can also select a custom algorithm, which is independent of the BIOS profile. The options available are:

- Maximum Performance (Performance Optimized) :
 - Reduced probability of memory or CPU throttling.
 - Increased probability of turbo mode activation.
 - Generally, higher fan speeds at idle and stress loads.
- Minimum Power (Performance per Watt Optimized):
 - Optimized for lowest system power consumption based on optimum fan power state.
 - Generally, lower fan speeds at idle and stress loads.
- **Sound Cap** Sound Cap provides reduced acoustical output from a server at the expense of some performance. Enabling Sound Cap may include temporary deployment or evaluation of a server in an occupied space, but it should not be used during benchmarking or performance sensitive applications.

(i) **NOTE:** Selecting **Maximum Performance** or **Minimum Power**, overrides thermal settings associated to System Profile setting under **System BIOS** > **System BIOS** Settings.System Profile Settings page.

• **Maximum Exhaust Temperature Limit** — From the drop-down menu, select the maximum exhaust air temperature. The values are displayed based on the system.

The default value is Default, 70°C (158 °F).

This option allows the system fans speeds to change such that the exhaust temperature does not exceed the selected exhaust temperature limit. This cannot always be guaranteed under all system operating conditions due to dependency on system load and system cooling capability.

• Fan Speed Offset — Selecting this option allows additional cooling to the server. In case hardware is added (example, new PCle cards), it may require additional cooling. A fan speed offset causes fan speeds to increase (by the offset % value) over baseline fan speeds calculated by the Thermal Control algorithm. Possible values are:



- Low Fan Speed Drives fan speeds to a moderate fan speed.
- Medium Fan Speed Drives fan speeds close to medium.
- **High Fan Speed** Drives fan speeds close to full speed.
- Max Fan Speed Drives fan speeds to full speed.
- **Off** Fan speed offset is set to off. This is the default value. When set to off, the percentage does not display. The default fan speed is applied with no offset. Conversely, the maximum setting will result in all fans running at maximum speed.

The fan speed offset is dynamic and based on the system. The fan speed increase for each offset is displayed next to each option.

The fan speed offset increases all fan speeds by the same percentage. Fan speeds may increase beyond the offset speeds based on individual component cooling needs. The overall system power consumption is expected to increase.

Fan speed offset allows you to increase the system fan speed with four incremental steps. These steps are equally divided between the typical baseline speed and the maximum speed of the server system fans. Some hardware configurations results in higher baseline fan speeds, which results in offsets other than the maximum offset to achieve maximum speed.

The most common usage scenario is non-standard PCIe adapter cooling. However, the feature can be used to increase system cooling for other purposes.

NOTE: FAN configuration setting is available in iDRAC even when system does not have any FANs. This is because, iDRAC sends the specified configuration to Chassis manager and Chassis manager can process the data from iDRAC and send the required cooling to the system as per configuration.

Thresholds

- **Maximum PCIe Inlet Temperature Limit** Default value is 55°C. Select the lower temperature of 45°C for third party PCIe cards which require lower inlet temperature.
 - **Exhaust Temperature Limits** By modifying the values for the following you can set the exhaust temperature limits:
 - Set Maximum Exhaust Temperature Limit
 - Set Air Temperature Rise Limit
- **Minimum Fan Speed in PWM (% of Max)** Select this option to fine tune the fan speed. Using this option, you can set a higher baseline system fan speed or increase the system fan speed if other custom fan speed options are not resulting in the required higher fan speeds.
 - **Default** Sets minimum fan speed to default value as determined by the system cooling algorithm.
 - **Custom** Enter the percentage by which you want to change the fan speed. Range is between 9-100.

The allowable range for minimum fan speed PWM is dynamic based on the system configuration. The first value is the idle speed and the second value is the configuration max (Depending on the system configuration, the maximum speed may be up to 100%.).

System fans can run higher than this speed as per thermal requirements of the system but not lower than the defined minimum speed. For example, setting Minimum Fan Speed at 35% limits the fan speed to never go lower than 35% PWM.

(i) NOTE: 0% PWM does not indicate fan is off. It is the lowest fan speed that the fan can achieve.

The settings are persistent, which means that once they are set and applied, they do not automatically change to the default setting during system reboot, power cycling, iDRAC, or BIOS updates. The custom cooling options may not be supported on all servers. If the options are not supported, they are not displayed or you cannot provide a custom value.

3. Click Apply to apply the settings.

The following message is displayed:

It is recommended to reboot the system when a thermal profile change has been made. This is to ensure all power and thermal settings are activated.

4. Click Reboot Later or Reboot Now.

(i) NOTE: You must reboot the system for the settings to take effect.



Modifying thermal settings using RACADM

To modify the thermal settings, use the objects in the **system.thermalsettings** group with the **set** sub command as provided in the following table.

Table 10. Thermal Settings

Object	Description	Usage	Example
AirExhaustTemp	irExhaustTempAllows you to set the maximum air exhaust temperature limit.Set to any of the following values (based on the system): 0 — Indicates 40°C0Indicates 40°C1 — Indicates 45°C02 — Indicates 50°C2 — Indicates 50°C	To check the existing setting on the system: racadm get system.thermalsetti ngs.AirExhaustTemp	
		 4 — Indicates 60°C 	The output is:
		• 255 — Indicates 70°C (default)	AirExhaustTemp=70
			This output indicates that the system is set to limit the air exhaust temperature to 70°C. To set the exhaust
			temperature limit to 60°C:
			racadm set system.thermalsetti ngs.AirExhaustTemp 4
			The output is:
		Object value modified successfully.	
		If a system does not support a particular air exhaust temperature limit, then when you run the following command:	
			racadm set system.thermalsetti ngs.AirExhaustTemp 0
			The following error message is displayed:
			ERROR: RAC947: Invalid object value specified.
			Make sure to specify the value depending on the type of object.
			For more information, see RACADM help.



Table 10. Thermal Settings (continued)

Object	Description	Usage	Example
			To set the limit to the default value:
			racadm set system.thermalsetti ngs.AirExhaustTemp 255
FanSpeedHighOffsetVal	 anSpeedHighOffsetVal Getting this variable reads the fan speed offset value in %PWM for High Fan Speed Offset setting. This value depends on the system. Use FanSpeedOffset object to set this value using index value 1. 	Values from 0-100	racadm get system.thermalsetti ngs FanSpeedHighOffsetV al
			A numerical value, for example 66, is returned. This value indicates that when you use the following command, it applies a fan speed offset of High (66% PWM) over the baseline fan speed
			racadm set system.thermalsetti ngs FanSpeedOffset 1
FanSpeedLowOffsetVal	 Getting this variable reads the fan speed offset value in %PWM for Low Fan Speed Offset setting. This value depends on the system 	Values from 0-100	racadm get system.thermalsetti ngs FanSpeedLowOffsetVa l
 system. Use FanSpeedOffset object to set this value using index value 0. 		This returns a value such as "23". This means that when you use the following command, it applies a fan speed offset of Low (23% PWM) over baseline fan speed	
			racadm set system.thermalsetti ngs FanSpeedOffset 0
FanSpeedMaxOffsetVal	 Getting this variable reads the fan speed offset value in %PWM for Max Fan Speed Offset setting. This value depends on the system 	Values from 0-100	racadm get system.thermalsetti ngs FanSpeedMaxOffsetVa l
	 Use FanSpeedOffset to set this value using index value 3 		This returns a value such as "100". This means that when you use the following command, it applies a fan speed offset of Max (meaning full speed, 100% PWM). Usually, this offset results in



Table 10. Thermal Settings (continued)

Object	Description	Usage	Example
			fan speed increasing to full speed.
			racadm set system.thermalsetti ngs FanSpeedOffset 3
FanSpeedMediumOffsetV	• Getting this variable reads	Values from 0-100	
al	the fan speed offset value in %PWM for Medium Fan Speed Offset setting.This value depends on the system.		racadm get system.thermalsetti ngs FanSpeedMediumOffse tVal
 Use FanSpeedOffset object to set this value using index value 2 		This returns a value such as "47". This means that when you use the following command, it applies a fan speed offset of Medium (47% PWM) over baseline fan speed	
			racadm set system.thermalsetti ngs FanSpeedOffset 2
FanSpeedOffset	edOffset • Using this object with	Values are: • 0 — Low Fan Speed • 1 — High Fan Speed • 2 — Medium Fan Speed • 3 — Max Fan Speed • 255 — None	To view the existing setting:
 Using this object with set command allows setting the required fan speed offset value. The index value decides 	existing Fan Speed Offset value.Using this object with set command allows sotting		racadm get system.thermalsetti ngs.FanSpeedOffset
	 The index value decides 		To set the fan speed offset to High value (as defined in FanSpeedHighOffsetVal)
	the offset that is applied and the FanSpeedLowOffsetVa 1, FanSpeedMaxOffsetVa		racadm set system.thermalsetti ngs.FanSpeedOffset 1
	l, FanSpeedHighOffsetV al, and FanSpeedMediumOffse tVal objects (defined earlier) are the values at which the offsets are applied.		
MFSMaximumLimit	Read Maximum limit for MFS	Values from 1 — 100	To display the highest value that can be set using MinimumFanSpeed option:
			racadm get system.thermalsetti ngs.MFSMaximumLimit



Table 10. Thermal Settings (continued)

Object	Description	Usage	Example
MFSMinimumLimit	Read Minimum limit for MFS	Values from 0 to MFSMaximumLimit Default is 255 (means None)	To display the lowest value that can be set using MinimumFanSpeed option. racadm get system.thermalsetti ngs.MFSMinimumLimit
MinimumFanSpeed	 Allows configuring the Minimum Fan speed that is required for the system to operate. It defines the baseline (floor) value for fan speed and system allows fans to go lower than this defined fan speed value. This value is %PWM value for fan speed. 	Values from MFSMinimumLimit to MFSMaximumLimit When get command reports 255, it means user configured offset is not applied.	To make sure that the system minimum speed does not decrease lower than 45% PWM (45 must be a value between MFSMinimumLimit to MFSMaximumLimit): racadm set system.thermalsetti ngs.MinimumFanSpeed 45
ThermalProfile	 Allows you to specify the Thermal Base Algorithm. Allows you to set the system profile as required for thermal behavior associated to the profile. 	 Values: 0 — Auto 1 — Maximum performance 2 — Minimum Power 	To view the existing thermal profile setting: racadm get system.thermalsetti ngs.ThermalProfile To set the thermal profile to Maximum Performance: racadm set system.thermalsetti ngs.ThermalProfile 1
ThirdPartyPCIFanRespo nse	 Thermal overrides for third-party PCI cards. Allows you to disable or enable the default system fan response for detected third-party PCI cards. You can confirm the presence of third-party PCI card by viewing the message ID PCI3018 in the Lifecycle Controller log. 	Values: • 1 — Enabled • 0 — Disabled (i) NOTE: The default value is 1.	To disable any default fan speed response set for a detected third-party PCI card: racadm set system.thermalsetti ngs.ThirdPartyPCIFa nResponse 0

Modifying thermal settings using iDRAC settings utility

To modify the thermal settings:

- In the iDRAC Settings utility, go to Thermal. The iDRAC Settings Thermal page is displayed.
- 2. Specify the following:
 - Thermal Profile
 - Maximum Exhaust Temperature Limit
 - Fan Speed Offset
 - Minimum Fan Speed

Setting up managed system 69



The settings are persistent, which means that once they are set and applied, they do not automatically change to the default setting during system reboot, power cycling, iDRAC, or BIOS updates. A few Dell servers may or may not support some or all of these custom user cooling options. If the options are not supported, they are not displayed or you cannot provide a custom value.

3. Click **Back**, click **Finish**, and then click **Yes**. The thermal settings are configured.

Modifying PCIe airflow settings using iDRAC web interface

Use the PCIe airflow settings when increased thermal margin is desired for custom high powered PCIe cards.

To modify the PCIe airflow settings:

- 1. In the iDRAC Web interface, go to Configuration > System Settings > Hardware Settings > Cooling Configuration. The PCIe Airflow Settings page is displayed below the fan settings section.
- 2. Specify the following:
 - LFM Mode Select the Custom mode to enable the custom LFM option.
 - **Custom LFM** Enter the LFM value.
- **3.** Click **Apply** to apply the settings.

The following message is displayed:

It is recommended to reboot the system when a thermal profile change has been made. This is to ensure all power and thermal settings are activated.

Click Reboot Later or Reboot Now.

(i) NOTE: You must reboot the system for the settings to take effect.

Setting up management station

A management station is a computer used for accessing iDRAC interfaces to remotely monitor and manage the PowerEdge server(s).

To set up the management station:

- 1. Install a supported operating system. For more information, see the release notes.
- 2. Install and configure a supported Web browser. For more information, see the release notes.
- **3.** Install the latest Java Runtime Environment (JRE) (required if Java plug-in type is used to access iDRAC using a Web browser).

(i) NOTE: You need Java 8 or later to use this feature and to launch iDRAC Virtual Console over an IPv6 network.

- 4. From the *Dell Systems Management Tools and Documentation* DVD, install Remote RACADM VMCLI from the SYSMGMT folder. Else, run **Setup** on the DVD to install Remote RACADM by default and other OpenManage software. For more information about RACADM, see *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.
- 5. Install the following based on the requirement:
 - SSH client
 - TFTP
 - Dell OpenManage Essentials

Accessing iDRAC remotely

To remotely access iDRAC Web interface from a management station, make sure that the management station is in the same network as iDRAC. For example:

• Blade servers — The management station must be on the same network as CMC and OME Modular. For more information on isolating CMC network from the managed system's network, see the *Chassis Management Controller User's Guide* available at https://www.dell.com/cmcmanuals.



• Rack and tower servers — Set the iDRAC NIC to Dedicated or LOM1 and make sure that the management station is on the same network as iDRAC.

To access the managed system's console from a management station, use Virtual Console through iDRAC Web interface.

Configuring supported web browsers

() NOTE: For information about the supported browsers and their versions, see the *Release Notes* available at https://www.dell.com/idracmanuals.

Most features of iDRAC web interface can be accessed using these browsers with default settings. For certain feature to work, you must change a few settings. These settings include disabling pop-up blockers, enabling Java, ActiveX, or HTML5 plug-in support and so on.

If you are connecting to iDRAC web interface from a management station that connects to the Internet through a proxy server, configure the web browser to access the Internet through this server.

NOTE: If you use Internet Explorer or Firefox to access the iDRAC web interface, you may need to configure certain settings as described in this section. You can use other supported browsers with their default settings.

(i) NOTE: Blank proxy settings are treated equivalent to No proxy.

Configuring Internet Explorer

This section provides details about configuring Internet Explorer (IE) to ensure you can access and use all features of the iDRAC web interface. These settings include:

- Resetting security settings
- Adding iDRAC IP to trusted sites
- Configuring IE to enable Active Directory SSO
- Disabling IE Enhanced Security Configuration

Resetting Internet Explorer security settings

Ensure that Internet Explorer (IE) settings are set to Microsoft-recommended defaults and customize the settings as described in this section.

- 1. Open IE as an administrator or using an administrator account.
- 2. Click Tools Internet Options Security Local Network or Local intranet.
- 3. Click Custom Level, select Medium-Low, and click Reset. Click OK to confirm.

Adding iDRAC IP to the trusted-sites list

When you access iDRAC Web interface, you are prompted to add iDRAC IP address to the list of trusted domains if the IP address is missing from the list. When completed, click **Refresh** or relaunch the web browser to establish a connection to iDRAC web interface. If you are not prompted to add the IP, it is recommended that you add the IP manually to the trusted-sites list.

() NOTE: When connecting to the iDRAC web interface with a certificate the browser does not trust, the browser's certificate error warning may display a second time after you acknowledge the first warning.

To add iDRAC IP address to the trusted-sites list:

- 1. Click Tools > Internet Options > Security > Trusted sites > Sites.
- 2. Enter the iDRAC IP address to the Add this website to the zone.
- 3. Click Add, click OK, and then click Close.
- 4. Click **OK** and then refresh your browser.



Configuring Internet Explorer to enable Active Directory SSO

To configure the browser settings for Internet Explorer:

- 1. In Internet Explorer, navigate to Local Intranet and click Sites.
- **2.** Select the following options only:
- Include all local (intranet) sites not listed on other zones.
 - Include all sites that bypass the proxy server.
- 3. Click Advanced.
- Add all relative domain names that will be used for iDRAC instances that is part of the SSO configuration (for example, myhost.example.com.)
- 5. Click Close and click OK twice.

Disabling Internet Explorer Enhanced Security Configuration

To ensure that you can download log files and other local elements using the web interface, it is recommended to disable Internet Explorer Enhanced Security Configuration from Windows features. For information about disabling this feature on your version of Windows, see Microsoft's documentation.

Configuring Mozilla Firefox

This section provides details about configuring Firefox to ensure you can access and use all features of the iDRAC web interface. These settings include:

- Disabling whitelist feature
- Configuring Firefox to enable Active Directory SSO

(i) NOTE: Mozilla Firefox browser may not have scroll bar for iDRAC Online Help page.

Disabling whitelist feature in Firefox

Firefox has a "whitelist" security feature that requires user permission to install plug-ins for each distinct site that hosts a plug-in. If enabled, the whitelist feature requires you to install a Virtual Console viewer for each iDRAC you visit, even though the viewer versions are identical.

To disable the whitelist feature and avoid unnecessary plug-in installations, perform the following steps:

- 1. Open a Firefox Web browser window.
- $\label{eq:linear} \textbf{2. In the address field, enter about: config and press < Enter>}.$
- 3. In the **Preference Name** column, locate and double-click **xpinstall.whitelist.required**.
- The values for **Preference Name**, **Status**, **Type**, and **Value** change to bold text. The **Status** value changes to user set and the **Value** changes to false.
- In the Preferences Name column, locate xpinstall.enabled.
 Make sure that Value is true. If not, double-click xpinstall.enabled to set Value to true.

Configuring Firefox to enable Active Directory SSO

To configure the browser settings for Firefox:

- 1. In Firefox address bar, enter about:config.
- 2. In Filter, enter network.negotiate.
- 3. Add the domain name to network.negotiate-auth.trusted-uris (using comma separated list.)
- 4. Add the domain name to network.negotiate-auth.delegation-uris (using comma separated list.)

Configuring web browsers to use virtual console

To use Virtual Console on your management station:

72 Setting up managed system



 Make sure that a supported version of the browser (Internet Explorer (Windows), or Mozilla Firefox (Windows or Linux), Google Chrome, Safari) is installed.

For more information about the supported browser versions, see the *Release Notes* available at https://www.dell.com/ idracmanuals.

- 2. To use Internet Explorer, set IE to Run As Administrator.
- 3. Configure the Web browser to use ActiveX, Java, or HTML5 plug-in.

ActiveX viewer is supported only with Internet Explorer. HTML5 or a Java viewer is supported on any browser.

INOTE: You need Java 8 or later to use this feature and to launch iDRAC Virtual Console over an IPv6 network.

- 4. Import the root certificates on the managed system to avoid the pop-ups that prompt you to verify the certificates.
- 5. Install the compat-libstdc++-33-3.2.3-61 related package.

NOTE: On Windows, the compat-libstdc++-33-3.2.3-61 related package may be included in the .NET framework package or the operating system package.

6. If you are using MAC operating system, select the **Enable access for assistive devices** option in the **Universal Access** window.

For more information, see the MAC operating system documentation.

Configuring Internet Explorer to use HTML5-based plug-in

The HTML5 virtual console and virtual media APIs are created by using HTML5 technology. The following are the advantages of HTML5 technology:

- Installation is not required on the client workstation.
- Compatibility is based on browser and is not based on the operating system or installed components.
- Compatible with most of the desktops and mobile platforms.
- Quick deployment and the client is downloaded as part of a web page.

You must configure Internet Explorer (IE) settings before you launch and run HTML5 based virtual console and virtual media applications. To configure the browser settings:

- 1. Disable pop-up blocker. To do this, click **Tools** > **Internet Options** > **Privacy** and clear the **Turn on Pop-up Blocker** check-box.
- 2. Start the HTML5 virtual console using any of the following methods:
 - In IE, click Tools > Compatibility View Settings and clear the Display intranet sites in Compatibility View checkbox.
 - In IE using an IPv6 address, modify the IPv6 address as follows:

https://[fe80::d267:e5ff:fef4:2fe9]/ to https://fe80--d267-e5ff-fef4-2fe9.ipv6literal.net/

Direct HTML5 virtual console in IE using an IPv6 address, modify the IPv6 address as follows:

https://[fe80::d267:e5ff:fef4:2fe9]/console to https://fe80--d267-e5fffef4-2fe9.ipv6-literal.net/console

To display the Title Bar information in IE, go to Control Panel > Appearance and Personalization > Personalization > Window Classic

Configuring Microsoft Edge to use HTML5-based plug-in

You must configure Edge settings before you launch and run HTML5 based virtual console and virtual media applications. To configure the browser settings:

- 1. Click Settings > View Advanced Settings and disable the Block pop-ups option.
- **2.** Modify the IPv6 address as follows :

```
https://2607:f2b1:f083:147::leb.ipv6:literal.net/restgui to https://2607-f2b1-
f083-147--leb.ipv6-literal.net/restgui
```



Configuring the web browser to use Java plug-in

Install a Java Runtime Environment (JRE) if you are using Firefox or IE and want to use the Java Viewer.

NOTE: Install a 32-bit or 64-bit JRE version on a 64-bit operating system or a 32-bit JRE version on a 32-bit operating system.

To configure IE to use Java plug-in:

- Disable automatic prompting for file downloads in Internet Explorer.
- Disable Enhanced Security Mode in Internet Explorer.

Configuring IE to use ActiveX plug-in

You must configure the IE browser settings before you start and run ActiveX based Virtual Console and Virtual Media applications. The ActiveX applications are delivered as signed CAB files from the iDRAC server. If the plug-in type is set to Native-ActiveX type in Virtual console, when you try to start the Virtual Console, the CAB file is downloaded to the client system and ActiveX based Virtual Console is started. Internet Explorer requires some configurations to download, install, and run these ActiveX based applications.

On 64-bit operating systems, you can install both 32-bit or 64-bit versions of Internet Explorer. You may use either of 32-bit or 64-bit, however, you must install the corresponding plug-in. For example, if you install the plug-in in the 64-bit browser and then open the viewer in a 32-bit browser, you must install the plug-in again.

(i) NOTE: You can use ActiveX plug-in only with Internet Explorer.

NOTE: To use ActiveX plug-in on systems with Internet Explorer 9, before configuring Internet Explorer, ensure that you disable the Enhanced Security Mode in Internet Explorer or in the server manager in Windows Server operating systems.

For ActiveX applications in Windows 7, Windows 2008, and Windows 10 configure the following Internet Explorer settings to use the ActiveX plug-in:

- 1. Clear the browser's cache.
- 2. Add iDRAC IP or host name to the Local Internet site list.
- 3. Reset the custom settings to **Medium-low** or change the settings to allow installation of signed ActiveX plug-ins.
- 4. Enable the browser to download encrypted content and to enable third-party browser extensions. To do this, go to Tools > Internet Options > Advanced, clear the Do not save encrypted pages to disk option, and select the Enable third-party browser extensions option.

(i) NOTE: Restart Internet Explorer for the Enable third-party browser extension setting to take effect.

- 5. Go to Tools > Internet Options > Security and select the zone in which you want to run the application.
- 6. Click Custom level. In the Security Settings window, do the following:
 - Select Enable for Automatic prompting for ActiveX controls.
 - Select Prompt for Download signed ActiveX controls.
 - Select Enable or Prompt for Run ActiveX controls and plugins.
 - Select Enable or Prompt for Script ActiveX controls marked safe for scripting.
- 7. Click OK to close the Security Settings window.
- 8. Click OK to close the Internet Options window.
 - **NOTE:** On systems with Internet Explorer 11, ensure that you add the iDRAC IP by clicking **Tools** > **Compatibility View settings**.
 - () NOTE:
 - The varying versions of Internet Explorer share **Internet Options**. Therefore, after you add the server to the list of *trusted sites* for one browser the other browser uses the same setting.
 - Before installing the ActiveX control, Internet Explorer may display a security warning. To complete the ActiveX control installation procedure, accept the ActiveX control when Internet Explorer prompts you with a security warning.



 If you get the error Unknown Publisher while launching Virtual Console, it may be caused because of the change to the code-signing certificate path. To resolve this error, you must download an addition key. Use a search engine to search for Symantec SO16958 and, from the search results, follow the instructions on the Symantec website.

Additional settings for Windows Vista or newer Microsoft operating systems

The Internet Explorer browsers in Windows Vista or newer operating systems have an additional security feature called *Protected Mode*.

To launch and run ActiveX applications in Internet Explorer browsers with Protected Mode:

- 1. Run IE as an administrator.
- 2. Go to Tools > Internet Options > Security > Trusted Sites.
- Make sure that the Enable Protected Mode option is not selected for Trusted Sites zone. Alternatively, you can add the iDRAC address to sites in the Intranet zone. By default, protected mode is turned off for sites in Intranet Zone and Trusted Sites zone.
- 4. Click Sites.
- 5. In the Add this website to the zone field, add the address of your iDRAC and click Add.
- 6. Click **Close** and then click **OK**.
- 7. Close and restart the browser for the settings to take effect.

Clearing browser cache

If you have issues when operating the Virtual Console, (out of range errors, synchronization issues, and so on) clear the browser's cache to remove or delete any old versions of the viewer that may be stored on the system and try again.

Clearing earlier Java versions

To clear older versions of Java viewer in Windows or Linux, do the following:

- At the command prompt, run javaws-viewer or javaws-uninstall. The Java Cache viewer is displayed.
- 2. Delete the items titled iDRAC Virtual Console Client.

Importing CA certificates to management station

When you launch Virtual Console or Virtual Media, prompts are displayed to verify the certificates. If you have custom Web server certificates, you can avoid these prompts by importing the CA certificates to the Java or ActiveX trusted certificate store.

For more information about Automatic Certificate Enrollment (ACE), see section Automatic Certificate Enrollment

Importing CA certificate to Java trusted certificate store

To import the CA certificate to the Java trusted certificate store:

- 1. Launch the Java Control Panel.
- 2. Click Security tab and then click Certificates.
- The **Certificates** dialog box is displayed.
- 3. From the Certificate type drop-down menu, select **Trusted Certificates**.
- **4.** Click **Import**, browse, select the CA certificate (in Base64 encoded format), and click **Open**. The selected certificate is imported to the Web start trusted certificate store.
- 5. Click Close and then click OK. The Java Control Panel window closes.



Importing CA certificate to ActiveX trusted certificate store

You must use the OpenSSL command line tool to create the certificate Hash using Secure Hash Algorithm (SHA). It is recommended to use OpenSSL tool 1.0.x and later since it uses SHA by default. The CA certificate must be in Base64 encoded PEM format. This is one-time process to import each CA certificate.

To import the CA certificate to the ActiveX trusted certificate store:

- 1. Open the OpenSSL command prompt.
- 2. Run a 8 byte hash on the CA certificate that is currently in-use on the management station using the command: openssl x509 -in (name of CA cert) -noout -hash

An output file is generated. For example, if the CA certificate file name is **cacert.pem**, the command is:

openssl x509 -in cacert.pem -noout -hash

The output similar to "431db322" is generated.

- 3. Rename the CA file to the output file name and include a ".0" extension. For example, 431db322.0.
- 4. Copy the renamed CA certificate to your home directory. For example, C:\Documents and Settings\<user> directory.

Viewing localized versions of web interface

iDRAC web interface is supported in the following languages:

- English (en-us)
- French (fr)
- German (de)
- Spanish (es)
- Japanese (ja)
- Simplified Chinese (zh-cn)

The ISO identifiers in parentheses denote the supported language variants. For some supported languages, resizing the browser window to 1024 pixels wide is required to view all features.

iDRAC Web interface is designed to work with localized keyboards for the supported language variants. Some features of iDRAC Web interface, such as Virtual Console, may require additional steps to access certain functions or letters. Other keyboards are not supported and may cause unexpected problems.

NOTE: See the browser documentation on how to configure or setup different languages and view localized versions of iDRAC Web interface.

Updating device firmware

Using iDRAC, you can update the iDRAC, BIOS, and all device firmware that is supported by using Lifecycle Controller update such as:

- Fibre Channel (FC) cards
- Diagnostics
- Operating System Driver Pack
- Network Interface Card (NIC)
- RAID Controller
- Power Supply Unit (PSU)
- NVMe PCIe devices
- SAS/SATA hard drives
- Backplane update for internal and external enclosures
- OS Collector

CAUTION: The PSU firmware update may take several minutes depending on the system configuration and PSU model. To avoid damaging the PSU, do not interrupt the update process or power on the system during PSU firmware update.



NOTE: When updating the PSU firmware for PowerEdge C series servers, ensure that all servers in the same chassis are powered OFF first. If any of the other servers in the chassis are powered ON, the update process fails.

You must upload the required firmware to iDRAC. After the upload is complete, the current version of the firmware installed on the device and the version being applied is displayed. If the firmware being uploaded is not valid, an error message is displayed. Updates that do not require a reboot are applied immediately. Updates that require a system reboot are staged and committed to run on the next system reboot. Only one system reboot is required to perform all updates.

() NOTE:

- When SEKM mode is enabled on a controller, iDRAC Firmware downgrade/upgrade shall fail when tried from a SEKM to a non-SEKM iDRAC version. iDRAC Firmware upgrade/downgrade shall pass when done within the SEKM versions.
- PERC firmware downgrade shall fail when SEKM is enabled.

After the firmware is updated, the **System Inventory** page displays the updated firmware version and logs are recorded.

The supported firmware image file types are:

- .exe Windows-based Dell Update Package (DUP). You must have Control and Configure Privilege to use this image file type.
- .d9 Contains both iDRAC and Lifecycle Controller firmware

For files with .exe extension, you must have the System Control privilege. The Remote Firmware Update licensed feature and Lifecycle Controller must be enabled.

For files with .d9 extension, you must have the Configure privilege.

(i) NOTE: Ensure that all nodes in the system are powered off before updating the PSU firmware.

(i) **NOTE:** After upgrading the iDRAC firmware, you may notice a difference in the time stamp displayed in the Lifecycle Controller log. Time displayed in LC Log is different from NTP/Bios-Time for few logs during idrac reset.

You can perform firmware updates by using the following methods:

- Uploading a supported image type, one at a time, from a local system or network share.
- Connecting to an FTP, TFTP, HTTP or HTTPS site or a network repository that contains Windows DUPs and a corresponding catalog file.

You can create custom repositories by using the Dell Repository Manager. For more information, see *Dell Repository Manager Data Center User's Guide*. iDRAC can provide a difference report between the BIOS and firmware installed on the system and the updates available in the repository. All applicable updates contained in the repository are applied to the system. This feature is available with iDRAC Enterprise or Datacenter license.

(i) NOTE: HTTP/HTTPS only supports with either digest authentication or no authentication.

• Scheduling recurring automated firmware updates by using the catalog file and custom repository.

There are multiple tools and interfaces that can be used to update the iDRAC firmware. The following table is applicable only to iDRAC firmware. The table lists the supported interfaces, image-file types, and whether Lifecycle Controller must be in enabled state for the firmware to be updated.

Table 11. Image file types and dependencies

	.D9	.D9 Image		RAC DUP
Interface	Supported	Requires LC enabled	Supported	Requires LC enabled
BMCFW64.exe utility	Yes	No	No	N/A
Racadm FWUpdate (old)	Yes	No	No	NZA
Racadm Update (new)	Yes	Yes	Yes	Yes
iDRAC UI	Yes	Yes	Yes	Yes
WSMan	Yes	Yes	Yes	Yes
In-band OS DUP	No	N/A	Yes	No
Redfish	Yes	N/A	Yes	N/A



The following table provides information on whether a system restart is required when firmware is updated for a particular component:

() NOTE: When multiple firmware updates are applied through out-of-band methods, the updates are ordered in the most efficient possible manner to reduce unnecessary system restart.

Table 12. Firmware update — supported components

Component Name	Firmware Rollback Supported? (Yes or No)	Out-of-band — System Restart Required?	In-band — System Restart Required?	Lifecycle Controller GUI — Restart Required?	
Diagnostics	No	No	No	No	
OS Driver Pack	No	No	No	No	
idrac	Yes	No	No*	Yes	
BIOS	Yes	Yes	Yes	Yes	
RAID Controller	Yes	Yes	Yes	Yes	
BOSS	Yes	Yes	Yes	Yes	
NVDIMM	No	Yes	Yes	Yes	
Backplanes	Yes	Yes	Yes	Yes	
 i) NOTE: • For Expander (Act • For SEP (Passive) 	ive) backplanes, system re backplanes, rebootless up	estart is required. date is supported only fror	n 4.00.00.00 release onwa	rds.	
Enclosures	Yes	Yes	No	Yes	
NIC	Yes	Yes	Yes	Yes	
Power Supply Unit	Yes	Yes	Yes	Yes	
CPLD	No	Yes	Yes	Yes	
(i) NOTE: After CPLD fir	mware upgrade is complet	e, iDRAC restarts automat	tically.		
FC Cards	Yes	Yes	Yes	Yes	
NVMe PCIe SSD drives	Yes	Yes	Yes	Yes	
() NOTE: Rebootless up	date is supported on some	e devices starting from rele	ease 5.00.00.00.		
SAS/SATA hard drives	No	Yes	Yes	No	
OS Collector	No	No	No	No	
CMC (on PowerEdge FX2 servers)	No	Yes	Yes	Yes	
ТРМ	No	Yes	Yes	Yes	
NOTE: TPM is supported from 5.00.00.00 release onwards and the action is staged. Only firmware update is supported. Downgrading and reinstalling the same firmware is not supported.					

(i) NOTE: For details of supported components for MX platform, see Table 13.

Table 13. Firmware update — supported components for MX platforms

Component Name	Firmware Rollback Supported? (Yes or No)	Out-of-band — System Restart Required?	In-band — System Restart Required?	Lifecycle Controller GUI — Restart Required?
Diagnostics	No	No	No	No
OS Driver Pack	No	No	No	No

78 Setting up managed system



Table 13. Firmware update — supported components for MX platforms (continued)

Component Name	Firmware Rollback Supported? (Yes or No)	Out-of-band — System Restart Required?	In-band — System Restart Required?	Lifecycle Controller GUI — Restart Required?
idrac	Yes	No	No*	Yes
BIOS	Yes	Yes	Yes	Yes
RAID Controller	Yes	Yes	Yes	Yes
BOSS	Yes	Yes	Yes	Yes
NVDIMM	No	Yes	Yes	Yes
Backplanes	Yes	Yes	Yes	Yes
Enclosures	Yes	Yes	No	Yes
NIC	Yes	Yes	Yes	Yes
Power Supply Unit	No	No	No	No
CPLD	No	Yes	Yes	Yes
FC Cards	Yes	Yes	Yes	Yes
NVMe PCle SSD drives	Yes	No	No	No
SAS/SATA hard drives	No	Yes	Yes	No
OS Collector	No	No	No	No

* Indicates that though a system restart is not required, iDRAC must be restarted to apply the updates. iDRAC communication and monitoring may temporarily be interrupted.

When you check for updates, the version marked as **Available** does not always indicate that it is the latest version available. Before you install the update, ensure that the version you choose to install is newer than the version currently installed. If you want to control the version that iDRAC detects, create a custom repository using Dell Repository Manager (DRM) and configure iDRAC to use that repository to check for updates.

Updating firmware using iDRAC web interface

You can update the device firmware using firmware images available on the local system, from a repository on a network share (CIFS, NFS, HTTP or HTTPS), or from FTP.

Updating single device firmware

Before updating the firmware using single device update method, make sure that you have downloaded the firmware image to a location on the local system.

(i) NOTE: Ensure that the file name for the single component DUP does not have any blank space.

To update single device firmware using iDRAC web interface:

1. Go to Maintenance > System Update.

The Firmware Update page is displayed.

2. On the Update tab, select Local as the Location Type.

() NOTE: If you select Local, ensure that you download the firmware image to a location on the local system. Select one file to be staged to iDRAC for update. You can select additional files one file at a time, for uploading to iDRAC. The files are uploaded to a temporary space on iDRAC and is limited to approximately 300MB.

- 3. Click Browse, select the firmware image file for the required component, and then click Upload.
- 4. After the upload is complete, the Update Details section displays each firmware file uploaded to iDRAC and its status.



If the firmware image file is valid and was successfully uploaded, the **Contents** column displays a plus icon (+) icon next to the firmware image file name. Expand the name to view the **Device Name**, **Current**, and **Available firmware version** information.

- 5. Select the required firmware file and do one of the following:
 - For firmware images that do not require a host system reboot, click **Install** (only available option). For example, iDRAC firmware file.
 - For firmware images that require a host system reboot, click Install and Reboot or Install Next Reboot.
 - To cancel the firmware update, click **Cancel**.

When you click **Install and Reboot**, or **Install Next Reboot**, the message Updating Job Queue is displayed.

To display the Job Queue page, click Job Queue. Use this page to view and manage the staged firmware updates or click OK to refresh the current page and view the status of the firmware update.

NOTE: If you navigate away from the page without saving the updates, an error message is displayed and all the uploaded content is lost.

() NOTE: You will not be able to proceed further, if the session gets expired after uploading the firmware file. This issue can be only resolved by RACADM reset.

() NOTE: After firmaware update is completed, an error message is displayed - RAC0508: An unexpected error

occurred. Wait for few minutes and retry the operation. If the problem persists, contact service provider. This is expected. You can wait for sometime and refresh the browser. Then you are re-directed to login page.

Scheduling automatic firmware updates

You can create a periodic recurring schedule for iDRAC to check for new firmware updates. At the scheduled date and time, iDRAC connects to the specified destination, checks for new updates, and applies or stages all applicable updates. A log file is created on the remote server, which contains information about server access and staged firmware updates.

It is recommended that you create a repository using Dell Repository Manager (DRM) and configure iDRAC to use this repository to check for and perform firmware updates. Using an internal repository enables you to control the firmware and versions available to iDRAC and helps avoid any unintended firmware changes.

(i) NOTE: For more information about DRM, see www.dell.com/openmanagemanuals > Repository Manager .

iDRAC Enterprise or Datacenter license is required to schedule automatic updates.

You can schedule automatic firmware updates using the iDRAC web interface or RACADM.

(i) NOTE: IPv6 address is not supported for scheduling automatic firmware updates.

Scheduling automatic firmware update using web interface

To schedule automatic firmware update using web Interface:

- () NOTE: Do not create the next scheduled occurrence of an automatic update job if a job is already Scheduled. It overwrites the current scheduled job.
- In the iDRAC web interface, go to Maintenance > System Update > Automatic Update. The Firmware Update page is displayed.
- 2. Click the Automatic Update tab.
- 3. Select the Enable Automatic Update option.
- 4. Select any of the following options to specify if a system reboot is required after the updates are staged:
 - Schedule Updates Stage the firmware updates but do not reboot the server.
 - Schedule Updates and reboot Server Enables server reboot after the firmware updates are staged.
- 5. Select any of the following to specify the location of the firmware images:
 - **Network** Use the catalog file from a network share (CIFS, NFS, HTTP or HTTPS, TFTP). Enter the network share location details.

NOTE: While specifying the network share settings, it is recommended to avoid special characters for user name and password or percent encode the special characters.



- **FTP** Use the catalog file from the FTP site. Enter the FTP site details.
- **HTTP** or **HTTPS** Allows catalog file streaming and via HTTP and HTTPS file transfer.

6. Based on the selection in step 5, enter the network settings or the FTP settings. For information about the fields, see the *iDRAC Online Help*.

7. In the **Update Window Schedule** section, specify the start time for the firmware update and the frequency of the updates (daily, weekly, or monthly).

For information about the fields, see the *iDRAC Online Help*.

8. Click Schedule Update.

The next scheduled job is created in the job queue. Five minutes after the first instance of the recurring job starts, the job for the next time period is created.

Scheduling automatic firmware update using RACADM

To schedule automatic firmware update, use the following commands:

• To enable automatic firmware update:

racadm set lifecycleController.lcattributes.AutoUpdate.Enable 1

• To view the status of automatic firmware update:

racadm get lifecycleController.lcattributes.AutoUpdate

• To schedule the start time and frequency of the firmware update:

```
racadm AutoUpdateScheduler create -u username -p password -l <location> [-f
catalogfilename -pu <proxyuser> -pp<proxypassword> -po <proxy port> -pt <proxytype>]
-time < hh:mm> [-dom < 1 - 28,L,'*'> -wom <l-4,L,'*'> -dow <sun-sat,'*'>] -rp <l-366>
-a <applyserverReboot (1-enabled | 0-disabled)>
```

For example,

• To automatically update firmware using a CIFS share:

racadm AutoUpdateScheduler create -u admin -p pwd -l //1.2.3.4/CIFS-share -f cat.xml -time 14:30 -wom 1 -dow sun -rp 5 -a 1

To automatically update firmware using FTP:

racadm AutoUpdateScheduler create -u admin -p pwd -l ftp.mytest.com -pu puser -pp puser -po 8080 -pt http -f cat.xml -time 14:30 -wom 1 -dow sun -rp 5 -a 1

• To view the current firmware update schedule:

racadm AutoUpdateScheduler view

• To disable automatic firmware update:

racadm set lifecycleController.lcattributes.AutoUpdate.Enable 0

• To clear the schedule details:

racadm AutoUpdateScheduler clear

Updating device firmware using RACADM

To update device firmware using RACADM, use the update subcommand. For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Examples:

• Upload the update file from a remote HTTP share:

```
racadm update -f <updatefile> -u admin -p mypass -l http://1.2.3.4/share
```



• Upload the update file from a remote HTTPS share:

```
racadm update -f <updatefile> -u admin -p mypass -l https://1.2.3.4/share
```

• To generate a comparison report using an update repository:

```
racadm update -f catalog.xml -l //192.168.1.1 -u test -p passwd --verifycatalog
```

 To perform all applicable updates from an update repository using myfile.xml as a catalog file and perform a graceful reboot:

racadm update -f "myfile.xml" -b "graceful" -l //192.168.1.1 -u test -p passwd

• To perform all applicable updates from an FTP update repository using Catalog.xml as a catalog file:

racadm update -f "Catalog.xml" -t FTP -e 192.168.1.20/Repository/Catalog

Updating firmware using CMC web interface

You can update iDRAC firmware for blade servers using the CMC Web interface.

To update iDRAC firmware using the CMC Web interface:

- 1. Log in to CMC Web interface.
- Go to iDRAC Settings > Settings > CMC. The Deploy iDRAC page is displayed.
- 3. Click Launch iDRAC Web interface and perform iDRAC Firmware Update.

Updating firmware using DUP

Before you update firmware using Dell Update Package (DUP), make sure to:

- Install and enable the IPMI and managed system drivers.
- Enable and start the Windows Management Instrumentation (WMI) service if your system is running Windows operating system,
 - **NOTE:** While updating the iDRAC firmware using the DUP utility in Linux, if you see error messages such as usb 5-2:
 - device descriptor read/64, error -71 displayed on the console, ignore them.
- If the system has ESX hypervisor installed, then for the DUP file to run, make sure that the "usbarbitrator" service is stopped using command: service usbarbitrator stop

Some versions of DUPs are constructed in ways that conflict with each other. This happens over time as new versions of the software are created. A newer version of software may drop support for legacy devices. Support for new devices may be added. Consider, for example, the two DUPs Network_Firmware_NDT09_WN64_21.60.5.EXE and Network_Firmware_8J1P7_WN64_21.60.27.50.EXE. The devices supported by these DUPs fit into three groups.

- Group A are legacy devices supported only by NDT09.
- Group B are devices supported by both NDT09 and 8J1P7.
- Group C are new devices supported only by 8J1P7.

Consider a server that has one or more devices from each of Groups A, B, and C. If the DUPs are used one at a time they should be successful. Using NDT09 by itself updates the devices in group A and group B. Using 8J1P7 by itself updates devices in group B and group C. However, if you try to use both DUPs at the same time that may attempt to create two updates for the Group B devices at the same time. That may fail with a valid error: "Job for this device is already present". The update software is unable to resolve the conflict of two valid DUPs attempting two valid updates on the same devices at the same time. At the same time both DUPs are required to support Group A and Group C devices. The conflict extends to performing rollbacks on the devices too. For best practice it is suggested to use each DUP individually.

To update iDRAC using DUP:

- 1. Download the DUP based on the installed operating system and run it on the managed system.
- 2. Run the DUP.
- The firmware is updated. A system restart is not required after firmware update is complete.



Updating firmware using remote RACADM

- 1. Download the firmware image to the TFTP or FTP server. For example, C:\downloads\firmimg.d9
- 2. Run the following RACADM command:

```
TFTP server:
```

• Using fwupdate command:

```
racadm -r <iDRAC IP address> -u <username> -p <password> fwupdate -g -u -a <path>
```

path

the location on the TFTP server where firmimg.d9 is stored.

• Using update command:

```
racadm -r <iDRAC IP address> -u <username> -p <password> update -f <filename>
```

FTP server:

• Using fwupdate command:

```
racadm -r <iDRAC IP address> -u <username> -p <password> fwupdate -f <ftpserver IP>
<ftpserver username> <ftpserver password> -d <path>
```

path

the location on the FTP server where firmimg.d9 is stored.

Using update command:

racadm -r <iDRAC IP address> -u <username> -p <password> update -f <filename>

For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Updating firmware using Lifecycle Controller Remote Services

For information to update the firmware using Lifecycle Controller–Remote Services, see *Lifecycle Controller Remote Services Quick Start Guide* available at https://www.dell.com/idracmanuals.

Updating CMC firmware from iDRAC

In the PowerEdge FX2/FX2s chassis, you can update the firmware for the Chassis Management Controller and any component that can be updated by CMC and shared by the servers from iDRAC.

Before applying the update, make sure that:

- Servers are not allowed to power-up by CMC.
- Chassis with LCD must display a message indicating "update is in-progress".
- Chassis without LCD must indicate the update progress using LED blinking pattern.
- During the update, chassis action power commands are disabled.

The updates for components such as Programmable System-on-Chip (PSoC) of IOM that requires all the servers to be idle, the update is applied on the next chassis power-up cycle.

CMC settings to update CMC firmware from iDRAC

In the PowerEdge FX2/FX2s chassis, before performing the firmware update from iDRAC for CMC and its shared components, do the following:

- 1. Launch the CMC Web interface
- 2. Go to iDRAC Settings > Settings > CMC. The Deploy iDRAC page is displayed.
- 3. From the Chassis Management at Server Mode , select Manage and Monitor, and the click Apply.



iDRAC settings to update CMC firmware

In the PowerEdge FX2/FX2s chassis, before updating the firmware for CMC and its shared components from iDRAC, do the following settings in iDRAC:

- 1. Go to iDRAC Settings > Settings > CMC.
- 2. Click on Chassis Management Controller Firmware Update The Chassis Management Controller Firmware Update Settings page is displayed.
- For Allow CMC Updates Through OS and Lifecycle Controller, select Enabled to enable CMC firmware update from iDRAC.
- 4. Under Current CMC Setting, make sure that Chassis Management at Server Mode option displays Manage and Monitor. You can set this in CMC.

Viewing and managing staged updates

You can view and delete the scheduled jobs including configuration and update jobs. This is a licensed feature. All jobs queued to run during the next reboot can be deleted.

Viewing and managing staged updates using iDRAC web interface

To view the list of scheduled jobs using iDRAC web interface, go to **Maintenance** > **Job Queue**. The **Job Queue** page displays the status of jobs in the Lifecycle Controller job queue. For information about the displayed fields, see the *iDRAC Online Help*.

To delete job(s), select the job(s) and click **Delete**. The page is refreshed and the selected job is removed from the Lifecycle Controller job queue. You can delete all the jobs queued to run during the next reboot. You cannot delete active jobs, that is, jobs with the status *Running* or *Downloading*.

You must have Server Control privilege to delete jobs.

Viewing and managing staged updates using RACADM

To view the staged updates using RACADM, use jobqueue sub-command. For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Rolling back device firmware

You can roll back the firmware for iDRAC or any device that Lifecycle Controller supports, even if the upgrade was previously performed using another interface. For example, if the firmware was upgraded using the Lifecycle Controller GUI, you can roll back the firmware using the iDRAC web interface. You can perform firmware rollback for multiple devices with one system reboot.

On Dell's 14th generation PowerEdge servers that have a single iDRAC and Lifecycle Controller firmware, rolling back the iDRAC firmware also rolls back the Lifecycle Controller firmware.

It is recommended to keep the firmware updated to ensure you have the latest features and security updates. You may need to rollback an update or install an earlier version if you encounter any issues after an update. To install an earlier version, use Lifecycle Controller to check for updates and select the version you want to install.

For the details about supported and unsupported components for firmware rollback, refer to the table Firmware update — supported components

You can perform firmware rollback for the following components:

- iDRAC with Lifecycle Controller
- BIOS
- Network Interface Card (NIC)
- Power Supply Unit (PSU)
- RAID Controller
- Backplane



(i) NOTE: You cannot perform firmware rollback for Diagnostics, Driver Packs, and CPLD.

Before rolling back the firmware, make sure that:

- You have Configure privilege to roll back iDRAC firmware.
- You have Server Control privilege and have enabled Lifecycle Controller to roll back firmware for any other device other than the iDRAC.
- Change the NIC mode to **Dedicated** if the mode is set as **Shared LOM**.

You can roll back the firmware to the previously installed version using any of the following methods:

- iDRAC web interface
- CMC web interface (not supported on MX platforms)
- OME-Modular web interface (Supported on MX platforms)
- CMC RACADM CLI (not supported on MX platforms)
- iDRAC RACADM CLI
- Lifecycle Controller GUI
- Lifecycle Controller-Remote Services

Rollback firmware using iDRAC web interface

To roll back device firmware:

1. In the iDRAC Web interface, go to Maintenance > System Update > Rollback.

The **Rollback** page displays the devices for which you can rollback the firmware. You can view the device name, associated devices, currently installed firmware version, and the available firmware rollback version.

- 2. Select one or more devices for which you want to rollback the firmware.
- 3. Based on the selected devices, click Install and Reboot or Install Next Reboot. If only iDRAC is selected, then click Install.
 - When you click Install and Reboot or Install Next Reboot, the message "Updating Job Queue" is displayed.

4. Click Job Queue.

The **Job Gueue** page is displayed, where you can view and manage the staged firmware updates.

() NOTE:

• While in rollback mode, the rollback process continues in the background even if you navigate away from this page.

An error message appears if:

- You do not have Server Control privilege to rollback any firmware other than the iDRAC or Configure privilege to rollback iDRAC firmware.
- Firmware rollback is already in-progress in another session.
- Updates are staged to run or already in running state.

If Lifecycle Controller is disabled or in recovery state and you try to perform a firmware rollback for any device other than iDRAC, an appropriate warning message is displayed along with steps to enable Lifecycle Controller.

Rollback firmware using CMC web interface

To roll back using the CMC Web interface:

- 1. Log in to CMC Web interface.
- Go to iDRAC Settings > Settings > CMC. The Deploy iDRAC page is displayed.
- **3.** Click **Launch iDRAC** and perform device firmware rollback as mentioned in the Rollback firmware using iDRAC web interface.

Rollback firmware using RACADM

1. Check the rollback status and the FQDD using the swinventory command:

racadm swinventory



For the device for which you want to rollback the firmware, the Rollback Version must be Available. Also, note the FQDD.

2. Rollback the device firmware using:

racadm rollback <FQDD>

For more information, see *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Rollback firmware using Lifecycle Controller

For information, see Lifecycle Controller User's Guide available at https://www.dell.com/idracmanuals .

Rollback firmware using Lifecycle Controller-Remote Services

For information, see Lifecycle Controller Remote Services Quick Start Guide available at https://www.dell.com/idracmanuals.

Recovering iDRAC

iDRAC supports two operating system images to make sure a bootable iDRAC. In the event of an unforeseen catastrophic error and you lose both boot paths:

- iDRAC bootloader detects that there is no bootable image.
- System Health and Identify LED is flashed at ~1/2 second rate. (LED is located on the back of a rack and tower servers and on the front of a blade server.)
- Bootloader is now polling the SD card slot.
- Format an SD card with FAT using a Windows operating system, or EXT3 using a Linux operating system.
- Copy **firmimg.d9** to the SD card.
- Insert the SD card into the server.
- Bootloader detects the SD card, turns the flashing LED to solid amber, reads the firmimg.d9, reprograms iDRAC, and then reboots iDRAC.

Monitoring iDRAC using other Systems Management tools

You can discover and monitor iDRAC using Dell Management Console or Dell OpenManage Essentials. You can also use Dell Remote Access Configuration Tool (DRACT) to discover iDRACs, update firmware, and set up Active Directory. For more information, see the respective user's guides.

Support Server Configuration Profile — Import and Export

Server Configuration Profile (SCP) allows you to import and export server configuration files.

(i) NOTE: You need admin privileges to perform Export and Import SCP task.

You can import and export from local management station, and from a Network Share via CIFS, NFS, HTTP or HTTPS. Using SCP, you can select and import or export component level configurations for BIOS, NIC and RAID. You can import and export SCP to the local management station or to a CIFS, NFS, HTTP, or HTTPS network share. You can either import and export individual profiles of iDRAC, BIOS, NIC, and RAID, or all of them together as a single file.

You can specify preview import or export of the SCP where the job is running and configuration result is generated but none of the configuration has applied.

A job is created once the import or export is initiated through the GUI. The status of the jobs can be viewed on the Job Queue page.

86 Setting up managed system



(i) NOTE: Only Host Name or IP Address are accepted for destination address.

- **NOTE:** You can browse to a specific location to import the server configuration files. You need to select the correct server configuration file that you want to import. For example, import.xml.
- **NOTE:** Depending on the exported file format (that you selected), the extension is added automatically. For example, export_system_config.xml.
- () NOTE: SCP applies the full configuration in a single job with minimal number of reboots. However, in a few system configurations some attributes change the operation mode of a device or may create subdevices with new attributes. When this occurs, SCP may be unable to apply all settings during a single job. Review the ConfigResult entries for the job to resolve any pending configuration settings.

SCP allows you to perform OS deployment (OSD) using a single xml/json file across multiple systems. You can also perform existing operations such as configurations and repository updates all at once.

SCP also allows to export and import of SSH public keys for all iDRAC users. There are 4 SSH public keys for all users.

Following are the steps for OS deployment using SCP:

- 1. Export SCP file
- 2. SCP file contains all the suppressed attributes that are needed to perform OSD.
- 3. Edit / update the OSD attributes and then perform import operation.
- 4. These OSD attributes are then validated by SCP orchestrator.
- 5. SCP orchestrator performs the configuration and repository updates specified in SCP file.
- 6. After configuration and updates are done, host OS shutdowns.

(i) NOTE: Only CIFS and NFS share is supported for hosting OS media.

- 7. SCP orchestrator initiates the OSD by attaching the drivers for the selected operating system and then initiates one time boot to the OS media present in NFS / Share.
- **8.** LCL shows the progress of the job.
- ${\bf 9.}~$ Once BIOS boots to the OS media, SCP job shows as Complete.
- 10. The attached media and OS media will be automatically detached after 65535 seconds or after the duration specified by OSD.1#ExposeDuration attribute.

Importing server configuration profile using iDRAC web interface

To import the server configuration profile:

- Go to Configuration > Server Configuration Profile The Server Configuration Profile page is displayed.
- 2. Select one of the following to specify the location type:
 - Local to import the configuration file saved in a local drive.
 - **Network Share** to import the configuration file from CIFS or NFS share.
 - HTTP or HTTPS to import the configuration file from a local file using HTTP/HTTPS file transfer.
 - **NOTE:** Depending on the location type, you must enter the Network Settings or HTTP/HTTPS settings. If proxy is configured for HTTP/HTTPS, proxy settings are also required.
- 3. Select the components listed in Import Components option.
- 4. Select the **Shutdown** type.
- 5. Select the **Maximum wait time** to specify the wait time before the system shuts down after the import is complete.
- 6. Click Import.

Exporting server configuration profile using iDRAC web interface

To export the server configuration profile:

- 1. Go to Configuration > Server Configuration Profile The Server Configuration Profile page is displayed.
- 2. Click Export.



- 3. Select one of the following to specify the location type:
 - Local to save the configuration file on a local drive.
 - **Network Share** to save the configuration file on a CIFS or NFS share.
 - **HTTP or HTTPS** to save the configuration file to a local file using HTTP/HTTPS file transfer.
 - () NOTE: Depending on the location type, you must enter the Network Settings or HTTP/HTTPS settings. If proxy is configured for HTTP/HTTPS, proxy settings are also required.
- 4. Select the components that you need to back up the configuration for.
- 5. Select the **Export type**, following are the options:
 - Basic
 - Replacement Export
 - Clone Export
- 6. Select an Export file format.
- 7. Select Additional export items.
- 8. Click Export.

Secure Boot Configuration from BIOS Settings or F2

UEFI Secure Boot is a technology that eliminates a major security void that may occur during a handoff between the UEFI firmware and UEFI operating system (OS). In UEFI Secure Boot, each component in the chain is validated and authorized against a specific certificate before it is allowed to load or run. Secure Boot removes the threat and provides software identity checking at every step of the boot—Platform firmware, Option Cards, and OS BootLoader.

The Unified Extensible Firmware Interface (UEFI) Forum—an industry body that develops standards for pre-boot software —defines Secure Boot in the UEFI specification. Computer system vendors, expansion card vendors, and operating system providers collaborate on this specification to promote interoperability. As a portion of the UEFI specification, Secure Boot represents an industry-wide standard for security in the pre-boot environment.

When enabled, UEFI Secure Boot prevents the unsigned UEFI device drivers from being loaded, displays an error message, and does not allow the device to function. You must disable Secure Boot to load the unsigned device drivers.

On the Dell 14th generation and later versions of PowerEdge servers, you can enable or disable the Secure Boot feature by using different interfaces (RACADM, WSMAN, REDFISH, and LC-UI).

Acceptable file formats

The Secure Boot policy contains only one key in PK, but multiple keys may reside in KEK. Ideally, either the platform manufacturer or platform owner maintains the private key corresponding to the public PK. Third parties (such as OS providers and device providers) maintain the private keys corresponding to the public keys in KEK. In this way, platform owners or third parties may add or remove entries in the db or dbx of a specific system.

The Secure Boot policy uses db and dbx to authorize pre-boot image file execution. For an image file to get executed, it must associate with a key or hash value in db, and not associate with a key or hash value in dbx. Any attempts to update the contents of db or dbx must be signed by a private PK or KEK. Any attempts to update the contents of PK or KEK must be signed by a private PK.

Table 14. Acceptable file formats

Policy Component	Acceptable File Formats	Acceptable File Extensions		Max records allowed
РК	X.509 Certificate (binary DER format only)	1.	.cer	One
		2.	.der	
		3.	.crt	
КЕК	X.509 Certificate (binary DER format only)	1.	.cer	More than one
	Public Key Store	2.	.der	
		3.	.crt	



Table 14. Acceptable file formats (continued)

Policy Component	Acceptable File Formats	Ac	ceptable File Extensions	Max records allowed
		4.	.pbk	
DB and DBX	Ad DBX X.509 Certificate (binary DER format only) EFI image (system BIOS will estavitate and impact image	1.	.cer	More than one
		2.	.der	
digest)	3.	.crt		
		4.	.efi	

The Secure Boot Settings feature can be accessed by clicking System Security under System BIOS Settings. To go to System BIOS Settings, press F2 when the company logo is displayed during POST.

- By default, Secure Boot is Disabled and the Secure Boot policy is set to Standard. To configure the Secure Boot Policy, you must enable Secure Boot.
- When the Secure Boot mode is set to Standard, it indicates that the system has default certificates and image digests or hash loaded from the factory. This caters to the security of standard firmware, drivers, option-roms, and boot loaders.
- To support a new driver or firmware on a server, the respective certificate must be enrolled into the DB of Secure Boot certificate store. Therefore, Secure Boot Policy must be configured to Custom.

When the Secure Boot Policy is configured as Custom, it inherits the standard certificates and image digests loaded in the system by default, which you can modify. Secure Boot Policy configured as Custom allows you to perform operations such as View, Export, Import, Delete, Delete All, Reset, and Reset All. Using these operations, you can configure the Secure Boot Policies.

Configuring the Secure Boot Policy to Custom enables the options to manage the certificate store by using various actions such as Export, Import, Delete, Delete All, Reset, and Rest All on PK, KEK, DB, and DBX. You can select the policy (PK / KEK / DB / DBX) on which you want to make the change and perform appropriate actions by clicking the respective link. Each section will have links to perform the Import, Export, Delete, and Reset operations. Links are enabled based on what is applicable, which depends on the configuration at the point of time. Delete All and Reset All are the operations that have impact on all the policies. Delete All deletes all the certificates and image digests in the Custom policy, and Rest All restores all the certificates and image digests from Standard or Default certificate store.

BIOS recovery

The BIOS recovery feature allows you to manually recover the BIOS from a stored image. The BIOS is checked when the system is powered-on and if a corrupt or compromised BIOS is detected, an error message is displayed. You can then initiate the process of BIOS recovery using RACADM. To perform a manual BIOS recovery, see the iDRAC RACADM Command Line Interface Reference Guide available at https://www.dell.com/idracmanuals.



Configuring iDRAC

iDRAC enables you to configure iDRAC properties, set up users, and set up alerts to perform remote management tasks.

Before you configure iDRAC, make sure that the iDRAC network settings and a supported browser is configured, and the required licenses are updated. For more information about the licensable feature in iDRAC, see iDRAC licenses .

You can configure iDRAC using:

- iDRAC Web Interface
- RACADM
- Remote Services (see Lifecycle Controller Remote Services User's Guide)
- IPMITool (see Baseboard Management Controller Management Utilities User's Guide)

To configure iDRAC:

- 1. Log in to iDRAC.
- **2.** Modify the network settings if required.

NOTE: If you have configured iDRAC network settings, using iDRAC Settings utility during iDRAC IP address setup, then ignore this step.

- 3. Configure interfaces to access iDRAC.
- 4. Configure front panel display.
- 5. Configure System Location if required.
- 6. Configure time zone and Network Time Protocol (NTP) if required.
- 7. Establish any of the following alternate communication methods to iDRAC:
 - IPMI or RAC serial
 - IPMI serial over LAN
 - IPMI over LAN
 - SSH
- 8. Obtain the required certificates.
- 9. Add and configure iDRAC users with privileges.
- 10. Configure and enable e-mail alerts, SNMP traps, or IPMI alerts.
- **11.** Set the power cap policy if required.
- **12.** Enable the Last Crash Screen.
- 13. Configure virtual console and virtual media if required.
- 14. Configure vFlash SD card if required.
- **15.** Set the first boot device if required.
- **16.** Set the OS to iDRAC Pass-through if required.

Topics:

- Viewing iDRAC information
- Modifying network settings
- Cipher suite selection
- FIPS mode
- Configuring services
- Using VNC client to manage remote server
- Configuring front panel display
- Configuring time zone and NTP
- Setting first boot device
- Enabling or disabling OS to iDRAC Pass-through
- Obtaining certificates
- Configuring multiple iDRACs using RACADM



Disabling access to modify iDRAC configuration settings on host system

Viewing iDRAC information

You can view the basic properties of iDRAC.

Viewing iDRAC information using web interface

In the iDRAC Web interface, go to **iDRAC Settings** > **Overview** to view the following information related to iDRAC. For information about the properties, see *iDRAC Online Help*.

iDRAC Details

- Device Type
- Hardware Version
- Firmware Version
- Firmware Update
- RAC time
- IPMI version
- Number of Possible Sessions
- Number of Current Sessions
- IPMI Version

iDRAC Service Module

Status

Connection View

- State
- Switch Connection ID
- Switch Port Connection ID

Current Network Settings

- iDRAC MAC Address
- Active NIC Interface
- DNS Domain Name

Current IPv4 Setting

- IPv4 Enabled
- DHCP
- Current IP Address
- Current Subnet Mask
- Current Gateway
- Use DHCP to Obtain DNS Server Address
- Current Preferred DNS Server
- Current Alternate DNS Server

Current IPv6 Settings

- IPv6 Enable
- Autoconfiguration
- Current IP Address
- Current IP Gateway
- Link Local Address
- Use DHCPv6 to obtain DNS
- Current Preferred DNS Server
- Current Alternate DNS Server



Viewing iDRAC information using RACADM

To view iDRAC information using RACADM, see getsysinfo or get sub-command details provided in the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Modifying network settings

After configuring the iDRAC network settings using the iDRAC Settings utility, you can also modify the settings through the iDRAC Web interface, RACADM, Lifecycle Controller, and Server Administrator (after booting to the operating system). For more information on the tools and privilege settings, see the respective user's guides.

To modify the network settings using iDRAC Web interface or RACADM, you must have **Configure** privileges.

(i) NOTE: Changing the network settings may terminate the current network connections to iDRAC.

Modifying network settings using web interface

To modify the iDRAC network settings:

- In the iDRAC Web interface, go to iDRAC Settings > Connectivity > Network > Network Settings. The Network page is displayed.
- 2. Specify the network settings, common settings, IPv4, IPv6, IPMI, and/or VLAN settings as per your requirement and click **Apply**.

If you select **Auto Dedicated NIC** under **Network Settings**, when the iDRAC has its NIC Selection as shared LOM (1, 2, 3, or 4) and a link is detected on the iDRAC dedicated NIC, the iDRAC changes its NIC selection to use the dedicated NIC. If no link is detected on the dedicated NIC, then the iDRAC uses the shared LOM. The switch from shared to dedicated time-out is five seconds and from dedicated to shared is 30 seconds. You can configure this time-out value using RACADM or WSMan.

For information about the various fields, see the *iDRAC Online Help*.

NOTE: If the iDRAC is using DHCP and has obtained a lease for its IP address, it is released back to the DHCP server's address pool when NIC or Ipv4 or DHCP is disabled.

Modifying network settings using local RACADM

To generate a list of available network properties, use the command

racadm get iDRAC.Nic

To use DHCP to obtain an IP address, use the following command to write the object DHCPEnable and enable this feature.

racadm set iDRAC.IPv4.DHCPEnable 1

The following example shows how the command may be used to configure the required LAN network properties:

racadm	set	iDRAC.Nic.Enable 1
racadm	set	iDRAC.IPv4.Address 192.168.0.120
racadm	set	iDRAC.IPv4.Netmask 255.255.0
racadm	set	iDRAC.IPv4.Gateway 192.168.0.120
racadm	set	iDRAC.IPv4.DHCPEnable 0
racadm	set	iDRAC.IPv4.DNSFromDHCP 0
racadm	set	iDRAC.IPv4.DNS1 192.168.0.5
racadm	set	iDRAC.IPv4.DNS2 192.168.0.6
racadm	set	iDRAC.Nic.DNSRegister 1
racadm	set	iDRAC.Nic.DNSRacName RAC-EK00002
racadm	set	iDRAC.Nic.DNSDomainFromDHCP 0
racadm	set	iDRAC.Nic.DNSDomainName MYDOMAIN

(i) NOTE: If iDRAC.Nic.Enable is set to **0**, the iDRAC LAN is disabled even if DHCP is enabled.

92 Configuring iDRAC


Configuring IP filtering

In addition to user authentication, use the following options to provide additional security while accessing iDRAC:

- IP filtering limits the IP address range of the clients accessing iDRAC. It compares the IP address of an incoming login to
 the specified range and allows iDRAC access only from a management station whose IP address is within the range. All other
 login requests are denied.
- When repeated login failures occur from a particular IP address, it prevents the address from logging in to iDRAC for a preselected time span. If you unsuccessfully log in up to two times, you are allowed to log in again only after 30 seconds. If you unsuccessfully log in more than two times, you are allowed to log in again only after 60 seconds.

(i) NOTE: This features supports upto 5 IP ranges. You van view / set this feature using RACADM and Redfish.

As login failures accumulate from a specific IP address, they are registered by an internal counter. When the user successfully logs in, the failure history is cleared and the internal counter is reset.

NOTE: When login attempts are prevented from the client IP address, few SSH clients may display the message: ssh exchange identification: Connection closed by remote host.

Configure IP filtering using iDRAC web interface

You must have Configure privilege to perform these steps.

To configure IP filtering:

1. In iDRAC Web interface, go to iDRAC Settings > Connectivity > Network > Network Settings > Advanced Network Settings.

The **Network** page is displayed.

- Click Advanced Network Settings. The Network Security page is displayed.
- **3.** Specify the IP filtering settings using **IP Range Address** and **IP Range Subnet Mask**. For more information about the options, see *iDRAC Online Help*.
- 4. Click Apply to save the settings.

Federal Information Processing Standards — FIPS is a set of standards used by the United States government agencies and contractors. FIPS Mode is intended to meet the requirements of FIPS 140-2 level 1. For more information about FIPS, refer to the FIPS User Guide for iDRAC, and CMC for non MX platforms.

(i) NOTE: Enabling FIPS Mode resets iDRAC to the default settings.

Configuring IP filtering using RACADM

You must have Configure privilege to perform these steps.

To configure IP filtering, use the following RACADM objects in the iDRAC.IPBlocking group:

- RangeEnable
- RangeAddr
- RangeMask

The RangeMask property is applied to both the incoming IP address and to the RangeAddr property. If the results are identical, the incoming login request is allowed to access iDRAC. Logging in from IP addresses outside this range results in an error.

(i) NOTE: Configuring IP filtering supports up to 5 IP ranges.

The login proceeds if the following expression equals zero:

```
RangeMask & (<incoming-IP-address> ^ RangeAddr)
```

&

Bitwise AND of the quantities

^



Bitwise exclusive-OR

Examples for IP Filtering

The following RACADM commands block all IP addresses except 192.168.0.57:

```
racadm set iDRAC.IPBlocking.RangeEnable 1
racadm set iDRAC.IPBlocking.RangeAddr 192.168.0.57
racadm set iDRAC.IPBlocking.RangeMask 255.255.255.255
```

To restrict logins to a set of four adjacent IP addresses (for example, 192.168.0.212 through 192.168.0.215), select all but the lowest two bits in the mask:

racadm set iDRAC.IPBlocking.RangeEnable 1
racadm set iDRAC.IPBlocking.RangeAddr 192.168.0.212
racadm set iDRAC.IPBlocking.RangeMask 255.255.255.252

The last byte of the range mask is set to 252, the decimal equivalent of 11111100b.

For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Cipher suite selection

Cipher Suite Selection can be used to limit the ciphers in iDRAC or client communications and determine how secure the connection will be. It provides another level of filtering the effective in-use TLS Cipher Suite. These settings can be configured through iDRAC web interface, RACADM, and WSMan command line interfaces.

Configuring cipher suite selection using iDRAC web interface

CAUTION: Using OpenSSL Cipher Command to parse strings with invalid syntax may lead to unexpected errors.

NOTE: This is an advanced security option. Before you configure this option, ensure that you have thorough knowledge of the following:

- The OpenSSL Cipher String Syntax and its use.
- Tools and Procedures to validate the resultant Cipher Suite Configuration to ensure that the results align with the
 expectations and requirements.
- **NOTE:** Before you configure the Advanced Settings for TLS Cipher Suites, ensure that you are using a supported web browser.

To add custom cipher strings:

- 1. In iDRAC web interface, go to iDRAC Settings > Services > Web Server.
- 2. Click Set Cipher String under the Customer Cipher String option.
- The Set Custom Cipher String page is displayed.
- 3. In the Custom Cipher String field, enter a valid string and click Set Cipher String.

(i) NOTE: For more information about cipher strings, see www.openssl.org/docs/man1.0.2/man1/ciphers.html.

4. Click Apply.

Setting the custom cipher string terminates the current iDRAC session. Wait for a few minutes before you open a new iDRAC session.

The ciphers supported by iDRAC on port 5000 are:

ssl-enum-ciphers:

TSLv1.1 Ciphers:

- TLS_ECDHE_RSA_WITH_3DES_EDE_CBC_SHA (secp256r1)
- TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA (secp256r1)
- TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA (secp256r1)
- TLS_ECDHE_RSA_WITH_RC4_128_SHA (secp256r1)
- TLS_RSA_WITH_3DES_EDE_CBC_SHA (rsa 2048)



- TLS_RSA_WITH_AES_128_CBC_SHA (rsa 2048)
- TLS_RSA_WITH_AES_256_CBC_SHA (rsa 2048)
- TLS_RSA_WITH_CAMELLIA_128_CBC_SHA (rsa 2048)
- TLS_RSA_WITH_CAMELLIA_256_CBC_SHA (rsa 2048)
- TLS_RSA_WITH_IDEA_CBC_SHA (rsa 2048)
- TLS_RSA_WITH_RC4_128_MD5 (rsa 2048)
- TLS_RSA_WITH_RC4_128_SHA (rsa 2048)
- TLS_RSA_WITH_SEED_CBC_SHA (rsa 2048)

TLSv1.2 Ciphers:

- TLS_ECDHE_RSA_WITH_3DES_EDE_CBC_SHA (secp256r1)
- TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA (secp256r1)
- TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256 (secp256r1)
- TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256 (secp256r1)
- TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA (secp256r1)
- TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384 (secp256r1)
- TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384 (secp256r1)
- TLS_ECDHE_RSA_WITH_RC4_128_SHA (secp256r1)
- TLS_RSA_WITH_3DES_EDE_CBC_SHA (rsa 2048)
- TLS_RSA_WITH_AES_128_CBC_SHA (rsa 2048)
- TLS_RSA_WITH_AES_128_CBC_SHA256 (rsa 2048)
- TLS_RSA_WITH_AES_128_GCM_SHA256 (rsa 2048)
- TLS_RSA_WITH_AES_256_CBC_SHA (rsa 2048)
- TLS_RSA_WITH_AES_256_CBC_SHA256 (rsa 2048)
- TLS_RSA_WITH_AES_256_GCM_SHA384 (rsa 2048)
- TLS_RSA_WITH_CAMELLIA_128_CBC_SHA (rsa 2048)
- TLS_RSA_WITH_CAMELLIA_256_CBC_SHA (rsa 2048)
- TLS_RSA_WITH_IDEA_CBC_SHA (rsa 2048)
- TLS_RSA_WITH_RC4_128_MD5 (rsa 2048)
- TLS_RSA_WITH_RC4_128_SHA (rsa 2048)
- TLS_RSA_WITH_SEED_CBC_SHA (rsa 2048)

Configuring cipher suite selection using RACADM

To configure cipher suite selection using RACADM, use any one of the following commands:

- racadm set idraC.webServer.customCipherString ALL:!DHE-RSA-AES256-GCM-SHA384:!DHE-RSA-AES256-GCM-SHA384
- racadm set idraC.webServer.customCipherString ALL:-DHE-RSA-CAMELLIA256-SHA
- racadm set idraC.webServer.customCipherString ALL:!DHE-RSA-AES256-GCM-SHA384:!DHE-RSA-
- AES256-SHA256:+AES256-GCM-SHA384:-DHE-RSA-CAMELLIA256-SHA

For more information about these objects, see *iDRAC RACADM Command Line Interface Reference Guide* available at dell.com/ idracmanuals.

FIPS mode

FIPS is a computer security standard that United States government agencies and contractors must use. Starting from version iDRAC 2.40.40.40, iDRAC supports enabling FIPS mode.

iDRAC will be officially certified to support FIPS mode in the future.

Difference between FIPS-mode supported and FIPS-validated

Software that has been validated by completing the Cryptographic Module Validation Program is referred to as FIPS-validated. Because of the time it takes to complete FIPS-validation, not all versions of iDRAC are validated. For information about the latest status of FIPS-validation for iDRAC, see the Cryptographic Module Validation Program page on the NIST website.



Enabling FIPS Mode

CAUTION: Enabling FIPS mode resets iDRAC to factory-default settings. If you want to restore the settings, back up the server configuration profile (SCP) before you enable FIPS mode, and restore the SCP after iDRAC restarts.

(i) NOTE: If you reinstall or upgrade iDRAC firmware, FIPS mode gets disabled.

Enabling FIPS mode using web interface

- 1. On the iDRAC web interface, navigate to iDRAC Settings > Connectivity > Network > Network Settings > Advanced Network Settings.
- 2. In FIPS Mode, select Enabled and click Apply.

(i) NOTE: Enabling FIPS Mode resets iDRAC to the default settings.

- **3.** A message appears prompting you to confirm the change. Click **OK**. iDRAC restarts in FIPS mode. Wait for at least 60 seconds before you reconnect to iDRAC.
- 4. Install a trusted certificate for iDRAC.

(i) NOTE: The default SSL certificate is not allowed in FIPS mode.

NOTE: Some iDRAC interfaces, such as the standards-compliant implementations of IPMI and SNMP, do not support FIPS-compliance.

Enabling FIPS mode using RACADM

Use RACADM CLI to execute the following command:

```
racadm set iDRAC.Security.FIPSMode <Enable>
```

Disabling FIPS mode

To disable FIPS mode, you must reset iDRAC to the factory-default settings.

Configuring services

You can configure and enable the following services on iDRAC:

Local Configuration	Disable access to iDRAC configuration (from the host system) using Local RACADM and iDRAC Settings utility.	
Web Server	Enable access to iDRAC web interface. If you disable the web interface, remote RACADM also gets disabled. Use local RACADM to re-enable the web server and remote RACADM.	
SEKM Configuration	Enables secure enterprise key management functionality on iDRAC using a client server architecture.	
SSH	Access iDRAC through firmware RACADM.	
Remote RACADM	Remotely access iDRAC.	
SNMP Agent	Enables support for SNMP queries (GET, GETNEXT, and GETBULK operations) in iDRAC.	
Automated System Recovery Agent	Enable Last System Crash Screen.	
Redfish	Enables support for Redfish RESTful API.	



Enable VNC server with or without SSL encryption.

Configuring services using web interface

To configure the services using iDRAC Web interface:

- In the iDRAC Web interface, go to iDRAC Settings > Services. The Services page is displayed.
- **2.** Specify the required information and click **Apply**.

For information about the various settings, see the *iDRAC Online Help*.

NOTE: Do not select the **Prevent this page from creating additional dialogs** check-box. Selecting this option prevents you from configuring services.

You can configure SEKM from iDRAC Settings page. Click iDRAC Settings > Services > SEKM Configuration.

(i) NOTE: For detailed step by step procedure for configuring SEKM, see the *iDRAC Online Help*.

NOTE: When **Security (Encryption)** mode is changed from **None** to **SEKM**, Real-Time job is not available. But it will be added to Staged job list. However, Real-Time job is successful when the mode is changed from **SEKM** to **None**.

Verify the following when changing the value of the **Username** Field in Client Certificate section on the KeySecure server (for ex: changing the value from **Common Name (CN)** to **User ID (UID)**)

- a. While using an existing account:
 - Verify in the iDRAC SSL certificate that, instead of the **Common Name** field, the **User name** field now matches the existing username on the KMS. If they don't, then you will have to set the username field and regenerate the SSL certificate again, get it signed on KMS and re-upload to iDRAC.
- **b.** While using a new user account:
 - Make sure the **User name** string matches the username field in the iDRAC SSL certificate.
 - If they don't match, then you will need to reconfigure the iDRAC KMS attributes Username and Password.
 - Once the certificate is verified to contain the username, then the only change that needs to be made is to change the key ownership from the old user to the new user to match the newly created KMS username.

While using Vormetric Data Security Manager as KMS, ensure that the Common Name (CN) field in iDRAC SSL certificate matches with the host name added to Vormetric Data Security Manager. Otherwise, the certificate may not import successfully.

() NOTE:

- Rekey option will be disabled when racadm sekm getstatus reports as Failed.
- SEKM only supports **Common name**, **User ID**, or **Organization Unit** for **User Name** field under Client certificate.
- If you are using a third party CA to sign the iDRAC CSR, ensure that the third party CA supports the value **UID** for **User Name** field in Client certificate. If it is not supported, use **Common Name** as the value for **User Name** field.
- If you are using Username and Password fields, ensure that KMS server supports those attributes.

NOTE: For KeySecure key management server,

- while creating an SSL certificate request, you must include the IP address of the key management server in **Subject** Alternative Name field
- the IP address must be in the following format: IP:xxx.xxx.xxx.xxx.

Configuring services using RACADM

To enable and configure services using RACADM, use the set command with the objects in the following object groups:

- iDRAC.LocalSecurity
- iDRAC.LocalSecurity
- iDRAC.SSH
- iDRAC.Webserver
- iDRAC.Racadm
- iDRAC.SNMP



For more information about these objects, see *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Enabling or disabling HTTPS redirection

If you do not want automatic redirection from HTTP to HTTPs due to certificate warning issue with default iDRAC certificate or as a temporary setting for debugging purpose, you can configure iDRAC such that redirection from http port (default is 80) to https port (default is 443) is disabled. By default, it is enabled. You have to log out and log in to iDRAC for this setting to take effect. When you disable this feature, a warning message is displayed.

You must have Configure iDRAC privilege to enable or disable HTTPS redirection.

An event is recorded in the Lifecycle Controller log file when this feature is enabled or disabled.

To disable the HTTP to HTTPS redirection:

racadm set iDRAC.Webserver.HttpsRedirection Disabled

To enable HTTP to HTTPS redirection:

racadm set iDRAC.Webserver.HttpsRedirection Enabled

To view the status of the HTTP to HTTPS redirection:

racadm get iDRAC.Webserver.HttpsRedirection

SEKM Functionalities

Following are the SEKM functionalities available in iDRAC:

- 1. SEKM Key Purge Policy iDRAC provides a policy setting that allows you to configure iDRAC to purge old unused keys at the Key Management Server (KMS) when Rekey operation is performed. You can set iDRAC read-writable attribute KMSKeyPurgePolicy to one of the following values:
 - Keep All Keys This is the default setting and is the existing behavior where iDRAC leaves all the keys on the KMS untouched while performing Rekey operation.
 - Keep N and N-1 keys iDRAC deletes all keys at the KMS except the current (N) and previous key (N-1) when performing Rekey operation.
- 2. KMS Key Purge on SEKM Disable As part of the Secure Enterprise Key Manager (SEKM) solution, iDRAC allows you to disable SEKM on the iDRAC. Once SEKM is disabled, the keys generated by iDRAC at the KMS are unused and remain at the KMS. This feature is for allowing iDRAC to delete those keys when SEKM is disabled. iDRAC provides a new option "-purgeKMSKeys" to existing legacy command "racadm sekm disable" which will let you purge keys at the KMS when SEKM is disabled on iDRAC.

NOTE: If SEKM is already disabled and you want to purge old keys, you must re-enable SEKM, then disable passing in option -purgeKMSKeys.

- **3. Key Creation Policy** As part of this release, iDRAC has been pre-configured with a Key Creation Policy. Attribute KeyCreationPolicy is read only and set to "Key per iDRAC" value.
- iDRAC read-only attribute iDRAC.SEKM.KeyIdentifierN reports the Key Identifier created by the KMS.

racadm get iDRAC.SEKM.KeyIdentifierN

• iDRAC read-only attribute iDRAC.SEKM.KeyldentifierNMinusOne reports the previous Key Identifier after performing a Rekey operation.

racadm get iDRAC.SEKM.KeyIdentifierNMinusOne

- 4. SEKM Rekey iDRAC provides 2 options to rekey your SEKM solution, either Rekey iDRAC or PERC. It's recommended to rekey the iDRAC since this rekeys all SEKM Secure capable/Enabled devices.
 - SEKM iDRAC Rekey [Rekey on iDRAC.Embedded.1 FQDD] When performing racadm sekm rekey iDRAC.Embedded.1, all SEKM Secure capable/Enabled devices are Rekeyed with a new key from KMS and this is common key to all SEKM enabled devices. iDRAC Rekey operation can also be executed from iDRAC GUI- **iDRAC Settings** > **Services** > **SEKM Configuration** > **Rekey**. After executing this operation, the change in the Key can be validated by reading KeyldentifierN and KeyldentifierNMinusOne attributes.



 SEKM PERC Rekey (Rekey On Controller [Example RAID.Slot.1-1] FQDD) — When performing racadm sekm rekey <controller FQDD>, the corresponding SEKM enabled controller gets rekeyed to the currently active iDRAC common key created from KMS. Storage Controller Rekey operation can also be executed from iDRAC GUI- Storage > Controllers > <controller FQDD> > Actions > Edit > Security > Security(Encryption) > Rekey.

Using VNC client to manage remote server

You can use a standard open VNC client to manage the remote server using both desktop and mobile devices such as Dell Wyse PocketCloud. When servers in data centers stop functioning, the iDRAC or the operating system sends an alert to the console on the management station. The console sends an email or SMS to a mobile device with required information and launches VNC viewer application on the management station. This VNC viewer can connect to OS/Hypervisor on the server and provide access to keyboard, video and mouse of the host server to perform the necessary remediation. Before launching the VNC client, you must enable the VNC server and configure the VNC server settings in iDRAC such as password, VNC port number, SSL encryption, and the time out value. You can configure these settings using iDRAC Web interface or RACADM.

(i) NOTE: VNC feature is licensed and is available in the iDRAC Enterprise or Datacenter license.

You can choose from many VNC applications or Desktop clients such as the ones from RealVNC or Dell Wyse PocketCloud.

2 VNC client sessions can be activated at the same time. Second one is in Read-Only mode.

If a VNC session is active, you can only launch the Virtual Media using Launch Virtual Console and not the Virtual Console Viewer.

If video encryption is disabled, the VNC client starts RFB handshake directly, and a SSL handshake is not required. During VNC client handshake (RFB or SSL), if another VNC session is active or if a Virtual Console session is open, the new VNC client session is rejected. After completion of the initial handshake, VNC server disables Virtual Console and allows only Virtual Media. After termination of the VNC session, VNC server restores the original state of Virtual Console (enabled or disabled).

() NOTE:

- While launching a VNC session, if you get an RFB protocol error, change the VNC client settings to High quality and then relaunch the session.
- When iDRAC NIC is in shared mode and the host system is power cycled, the network connection is lost for a few seconds. During this time, if you perform any action in the active VNC client, the VNC session may close. You must wait for timeout (value configured for the VNC Server settings in the **Services** page in iDRAC Web interface) and then re-establish the VNC connection.
- If the VNC client window is minimized for more than 60 seconds, the client window closes. You must open a new VNC session. If you maximize the VNC client window within 60 seconds, you can continue to use it.

Configuring VNC server using iDRAC web interface

To configure the VNC server settings:

- 1. In the iDRAC Web interface, go to **Configuration** > **Virtual Console**. The **Virtual Console** page is displayed.
- 2. In the VNC Server section, enable the VNC server, specify the password, port number, and enable or disable SSL encryption.

For information about the fields, see the *iDRAC Online Help*.

3. Click **Apply**. The VNC server is configured.

Configuring VNC server using RACADM

To configure the VNC server, use the set command with the objects in VNCserver.

For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.



Setting up VNC viewer with SSL encryption

While configuring the VNC server settings in iDRAC, if the **SSL Encryption** option was enabled, then the SSL tunnel application must be used along with the VNC Viewer to establish the SSL encrypted connection with iDRAC VNC server.

(i) NOTE: Most of the VNC clients do not have built-in SSL encryption support.

To configure the SSL tunnel application:

- 1. Configure SSL tunnel to accept connection on <localhost>:<localport number>. For example, 127.0.0.1:5930.
- Configure SSL tunnel to connect to <iDRAC IP address>:<VNC server port Number>. For example, 192.168.0.120:5901.
- 3. Start the tunnel application.

To establish connection with the iDRAC VNC server over the SSL encrypted channel, connect the VNC viewer to the localhost (link local IP address) and the local port number (127.0.0.1:<local port number>).

Setting up VNC viewer without SSL encryption

In general, all Remote Frame Buffer (RFB) compliant VNC Viewers connect to the VNC server using the iDRAC IP address and port number that is configured for the VNC server. If the SSL encryption option is disabled when configuring the VNC server settings in iDRAC, then to connect to the VNC Viewer do the following:

In the VNC Viewer dialog box, enter the iDRAC IP address and the VNC port number in the VNC Server field.

The format is <iDRAC IP address:VNC port number>

For example, if the iDRAC IP address is 192.168.0.120 and VNC port number is 5901, then enter 192.168.0.120:5901.

Configuring front panel display

You can configure the front panel LCD and LED display for the managed system.

- For rack and tower servers, two types of front panels are available:
- LCD front panel and System ID LED
- LED front panel and System ID LED

For blade servers, only the System ID LED is available on the server front panel since the blade chassis has the LCD.

Configuring LCD setting

You can set and display a default string such as iDRAC name, IP, and so on or a user-defined string on the LCD front panel of the managed system.

Configuring LCD setting using web interface

To configure the server LCD front panel display:

- 1. In iDRAC Web interface, go to Configurations > System Settings > Hardware Settings > Front Panel configuration.
- 2. In LCD Settings section, from the Set Home Message drop-down menu, select any of the following:
 - Service Tag (default)
 - Asset Tag
 - DRAC MAC Address
 - DRAC IPv4 Address
 - DRAC IPv6 Address
 - System Power
 - Ambient Temperature
 - System Model
 - Host Name
 - User Defined



None

If you select **User Defined**, enter the required message in the text box.

If you select **None**, home message is not displayed on the server LCD front panel.

- **3.** Enable Virtual Console indication (optional). If enabled, the Live Front Panel Feed section and the LCD panel on the server displays the Virtual console session active message when there is an active Virtual Console session.
- 4. Click Apply.

The server LCD front panel displays the configured home message.

Configuring LCD setting using RACADM

To configure the server LCD front panel display, use the objects in the System.LCD group.

For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Configuring LCD setting using iDRAC settings utility

To configure the server LCD front panel display:

- In the iDRAC Settings utility, go to Front Panel Security. The iDRAC Settings.Front Panel Security page is displayed.
- 2. Enable or disable the power button.
- **3.** Specify the following:
 - Access to the front panel
 - LCD message string
 - System power units, ambient temperature units, and error display
- 4. Enable or disable the virtual console indication.

For information about the options, see the *iDRAC Settings Utility Online Help*.

5. Click Back, click Finish, and then click Yes.

Configuring system ID LED setting

To identify a server, enable or disable System ID LED blinking on the managed system.

Configuring system ID LED setting using web interface

To configure the System ID LED display:

- 1. In iDRAC Web interface, go to Configuration > System Settings > Hardware Settings > Front Panel configuration. The System ID LED Settings page is displayed.
- 2. In System ID LED Settings section, select any of the following options to enable or disable LED blinking:
 - Blink Off
 - Blink On
 - Blink On 1 Day Timeout
 - Blink On 1 Week Timeout
 - Blink On 1 Month Timeout
- 3. Click Apply.

The LED blinking on the front panel is configured.

Configuring system ID LED setting using RACADM

To configure system ID LED, use the setled command.

For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.



Configuring time zone and NTP

You can configure the time zone on iDRAC and synchronize the iDRAC time using Network Time Protocol (NTP) instead of BIOS or host system times.

You must have Configure privilege to configure time zone or NTP settings.

Configuring time zone and NTP using iDRAC web interface

To configure time zone and NTP using iDRAC web interface:

- Go to iDRAC Settings > Settings > Time zone and NTP Settings. The Time zone and NTP page is displayed.
- 2. To configure the time zone, from the **Time Zone** drop-down menu, select the required time zone, and then click **Apply**.
- **3.** To configure NTP, enable NTP, enter the NTP server addresses, and then click **Apply**. For information about the fields, see *iDRAC Online Help*.

Configuring time zone and NTP using RACADM

To configure time zone and NTP, use the set command with the objects in the iDRAC.Time and iDRAC.NTPConfigGroup group.

For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

NOTE: iDRAC syncs the time with the host (local time). Hence it is recommended to configure both iDRAC and host with the same time zone so that the time sync is proper. If you want to change a time zone, you need to change it on both host and iDRAC and then the host needs to reboot.

Setting first boot device

You can set the first boot device for the next boot only or for all subsequent reboots. If you set the device to be used for all subsequent boots, it remains as the first boot device in the BIOS boot order until it is changed again either from the iDRAC web interface or from the BIOS boot sequence.

You can set the first boot device to one of the following:

- Normal Boot
- PXE
- BIOS Setup
- Local Floppy/Primary Removable Media
- Local CD/DVD
- Hard Drive
- Virtual Floppy
- Virtual CD/DVD/ISO
- Local SD Card
- Lifecycle Controller
- BIOS Boot Manager
- UEFI Device Path
- UEFI HTTP

() NOTE:

- BIOS Setup (F2), Lifecycle Controller (F10), and BIOS Boot Manager (F11) cannot be set as permanent boot device.
- The first boot device setting in iDRAC Web Interface overrides the System BIOS boot settings.

Setting first boot device using web interface

To set the first boot device using iDRAC Web interface:



- Go to Configuration > System Settings > Hardware Settings > First Boot Device. The First Boot Device page is displayed.
- Select the required first boot device from the drop-down list, and click Apply. The system boots from the selected device for subsequent reboots.
- To boot from the selected device only once on the next boot, select Boot Once. Thereafter, the system boots from the first boot device in the BIOS boot order.

For more information about the options, see the *iDRAC Online Help*.

Setting first boot device using RACADM

- To set the first boot device, use the iDRAC.ServerBoot.FirstBootDevice object.
- To enable boot once for a device, use the iDRAC.ServerBoot.BootOnce object.

For more information about these objects, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Setting first boot device using virtual console

You can select the device to boot from as the server is being viewed in the Virtual Console viewer before the server runs through its boot-up sequence. Boot-once is supported by all devices listed in Setting first boot device.

To set the first boot device using Virtual Console:

- 1. Launch Virtual Console.
- 2. In the Virtual Console Viewer, from the Next Boot menu, set the required device as the first boot device.

Enabling last crash screen

To troubleshoot the cause of a crash on the managed system, you can capture the system crash image using iDRAC. **NOTE:** For information about Server Administrator, see the *OpenManage Installation Guide* available at https://www.dell.com/openmanagemanuals.

The host system should have Windows Operating system to use this feature.

() NOTE:

- This feature is not applicable on Linux system.
- This feature is independent of any agents or attributes.

Enabling or disabling OS to iDRAC Pass-through

In servers that have Network Daughter Card (NDC) or embedded LAN On Motherboard (LOM) devices, you can enable the OS to iDRAC Pass-through feature. This feature provides a high-speed bi-directional in-band communication between iDRAC and the host operating system through a shared LOM, a dedicated NIC, or through the USB NIC. This feature is available for iDRAC Enterprise or Datacenter license.

NOTE: iDRAC Service Module (iSM) provides more features for managing iDRAC through the operating system. For more information, see the iDRAC Service Module User's Guide available at www.dell.com/idracservicemodule.

When enabled through dedicated NIC, you can launch the browser in the host operating system and then access the iDRAC Web interface. The dedicated NIC for the blade servers is through the Chassis Management Controller.

Switching between dedicated NIC or shared LOM does not require a reboot or reset of the host operating system or iDRAC.

You can enable this channel using:

- iDRAC web interface
- RACADM or WSMan (post operating system environment)
- iDRAC Settings utility (pre-operating system environment)

If the network configuration is changed through iDRAC Web interface, you must wait for at least 10 seconds before enabling OS to iDRAC Pass-through.



If you are configuring the server using a Server Configuration Profile through RACADM, WSMan or Redfish and if the network settings are changed in this file, then you must wait for 15 seconds to either enable OS to iDRAC Pass-through feature or set the OS Host IP address.

Before enabling OS to iDRAC Pass-through, make sure that:

- iDRAC is configured to use dedicated NIC or shared mode (that is, NIC selection is assigned to one of the LOMs).
- Host operating system and iDRAC are in the same subnet and same VLAN.
- Host operating system IP address is configured.
- A card that supports OS to iDRAC Pass-through capability is installed.
- You have the Configure privilege.

When you enable this feature:

- In shared mode, the host operating system's IP address is used.
- In dedicated mode, you must provide a valid IP address of the host operating system. If more than one LOM is active, enter the first LOM's IP address.

If the OS to iDRAC Pass-through feature does not work after it is enabled, ensure that you check the following:

- The iDRAC dedicated NIC cable is connected properly.
- At least one LOM is active.

NOTE: Use the default IP address. Ensure that the IP address of the USB NIC interface is not in the same network subnet as the iDRAC or host OS IP addresses. If this IP address conflicts with an IP address of other interfaces of the host system or the local network, you must change it.

- **NOTE:** If you launch iDRAC Service Module while USB NIC is in disabled state, the iDRAC Service Module changes the USB NIC IP address to 169.254.0.1.
- **NOTE:** Do not use 169.254.0.3 and 169.254.0.4 IP addresses. These IP addresses are reserved for the USB NIC port on the front panel when an A/A cable is used.
- () NOTE: iDRAC may not be accessible from the host server using LOM-Passthrough when NIC teaming is enabled. Then, iDRAC can be accessed from the host server OS using the iDRAC USB NIC or through the external network, via the iDRAC dedicated NIC.

Supported cards for OS to iDRAC Pass-through

The following table provides a list of cards that support the OS to iDRAC Pass-through feature using LOM.

Table 15. OS to iDRAC Pass-through using LOM — supported cards

Category	Manufacturer	Туре
NDC	Broadcom	• 5720 QP rNDC 1G BASE-T
	Intel	• x520/i350 QP rNDC 1G BASE-T

In-built LOM cards also support the OS to iDRAC pass-through feature.

Supported operating systems for USB NIC

The operating systems supported for USB NIC are:

- Server 2012 R2 Foundation Edition
- Server 2012 R2 Essentials Edition
- Server 2012 R2 Standard Edition
- Server 2012 R2 Datacenter Edition
- Server 2012 for Embedded Systems (Base and R2 w/ SP1)
- Server 2016 Essentials Edition
- Server 2016 Standard Edition
- Server 2016 Datacenter Edition
- RHEL 7.3
- RHEL 6.9



- SLES 12 SP2
- ESXi 6.0 U3
- vSphere 2016
- XenServer 7.1

For Linux operating systems, configure the USB NIC as DHCP on the host operating system before enabling USB NIC.

For vSphere, you must install the VIB file before enabling USB NIC.

NOTE: To configure USB NIC as DHCP in Linux operating system or XenServer, refer to the operating system or hypervisor documentation.

Installing VIB file

For vSphere operating systems, before enabling the USB NIC, you must install the VIB file.

To install the VIB file:

- 1. Using Win-SCP, copy the VIB file to /tmp/ folder of the ESX-i host operating system.
- 2. Go to the ESXi prompt and run the following command:

```
esxcli software vib install -v /tmp/ iDRAC USB NIC-1.0.0-799733X03.vib --no-sig-check
```

The output is:

```
Message: The update completed successfully, but the system needs to be rebooted for
the changes to be effective.
Reboot Required: true
VIBs Installed: Dell_bootbank_iDRAC_USB_NIC_1.0.0-799733X03
VIBs Removed:
VIBs Skipped:
```

- 3. Reboot the server.
- At the ESXi prompt, run the command: esxcfg-vmknic -1. The output displays the usb0 entry.

Enabling or disabling OS to iDRAC Pass-through using web interface

To enable OS to iDRAC Pass-through using Web interface:

- Go to iDRAC Settings > Connectivity > Network > OS to iDRAC Pass-through. The OS to iDRAC Pass-through page is displayed.
- 2. Change the State to **Enabled**.
- **3.** Select any of the following options for Pass-through Mode:
 - LOM The OS to iDRAC pass-through link between the iDRAC and the host operating system is established through the LOM or NDC.
 - **USB NIC** The OS to iDRAC pass-through link between the iDRAC and the host operating system is established through the internal USB bus.
 - (i) NOTE: If you set the pass-through mode to LOM, ensure that:
 - OS and iDRAC are on the same subnet
 - NIC selection in Network Settings is set to LOM
- 4. If the server is connected in shared LOM mode, then the OS IP Address field is disabled.

NOTE: If the VLAN is enabled on the iDRAC, the LOM-Passthrough will only function in shared LOM mode with VLAN tagging configured on the host.

() NOTE:

• When Pass-through mode is set to LOM, it is not possible to launch iDRAC from host OS after cold boot.



- We have purposefully removed the LOM Pass-through using Dedicated mode feature.
- 5. If you select **USB NIC** as the pass-through configuration, enter the IP address of the USB NIC.

The default value is 169.254.1.1. It is recommended to use the default IP address. However, if this IP address conflicts with an IP address of other interfaces of the host system or the local network, you must change it.

Do not enter 169.254.0.3 and 169.254.0.4 IPs. These IPs are reserved for the USB NIC port on the front panel when an A/A cable is used.

NOTE: If IPv6 is preferred, the default address is fde1:53ba:e9a0:de11::1. If needed, this address can be modified in the idrac.OS-BMC.UsbNicULA setting. If IPv6 is not wanted on the USB-NIC, it can be disabled by changing the address to "::"

- 6. Click Apply.
- 7. Click **Test Network Configuration** to check if the IP is accessible and the link is established between the iDRAC and the host operating system.

Enabling or disabling OS to iDRAC Pass-through using RACADM

To enable or disable OS to iDRAC Pass-through using RACADM, use the objects in the iDRAC.OS-BMC group. For more information, see theiDRAC Attribute Registry available at https://www.dell.com/idracmanuals.

Enabling or disabling OS to iDRAC Pass-through using iDRAC settings utility

To enable or disable OS to iDRAC Pass-through using iDRAC Settings Utility:

- In the iDRAC Settings utility, go to Communications Permissions. The iDRAC Settings.Communications Permissions page is displayed.
- 2. Select any of the following options to enable OS to iDRAC pass-through:
 - LOM The OS to iDRAC pass-through link between the iDRAC and the host operating system is established through the LOM or NDC.
 - **USB NIC** The OS to iDRAC pass-through link between the iDRAC and the host operating system is established through the internal USB bus.
 - **NOTE:** If you set the pass-through mode to LOM, ensure that:
 - OS and iDRAC are on the same subnet
 - NIC selection in Network Settings is set to a LOM
 - To disable this feature, select **Disabled**.

NOTE: The LOM option can be selected only of the card supports OS to iDRAC pass-through capability. Else, this option is grayed-out.

3. If you select **LOM** as the pass-through configuration, and if the server is connected using dedicated mode, enter the IPv4 address of the operating system.

i) NOTE: If the server is connected in shared LOM mode, then the OS IP Address field is disabled.

4. If you select USB NIC as the pass-through configuration, enter the IP address of the USB NIC.

The default value is 169.254.1.1. However, if this IP address conflicts with an IP address of other interfaces of the host system or the local network, you must change it. Do not enter 169.254.0.3 and 169.254.0.4 IPs. These IPs are reserved for the USB NIC port on the front panel when an A/A cable is used.

() NOTE: If IPv6 is preferred, the default address is fde1:53ba:e9a0:de11::1. If needed, this address can be modified in the idrac.OS-BMC.UsbNicULA setting. If IPv6 is not wanted on the USB-NIC, it can be disabled by changing the address to "::"

5. Click **Back**, click **Finish**, and then click **Yes**.

The details are saved.



Obtaining certificates

The following table lists the types of certificates based on the login type.

Table 16. Types of certificate based on login type

Login Type	Certificate Type	How to Obtain
Single Sign-on using Active Directory	Trusted CA certificate	Generate a CSR and get it signed from a Certificate Authority
		SHA-2 certificates are also supported.
Smart Card login as a local or Active Directory user	 User certificate Trusted CA certificate 	 User Certificate — Export the smart card user certificate as Base64-encoded file using the card management software provided by the smart card vendor. Trusted CA certificate — This certificate is issued by a CA. SHA-2 certificates are also supported.
Active Directory user login	Trusted CA certificate	This certificate is issued by a CA. SHA-2 certificates are also supported.
Local User login	SSL Certificate	 Generate a CSR and get it signed from a trusted CA i NOTE: iDRAC ships with a default self-signed SSL server certificate. The iDRAC Web server, Virtual Media, and Virtual Console use this certificate. SHA-2 certificates are also supported.

SSL server certificates

iDRAC includes a web server that is configured to use the industry-standard SSL security protocol to transfer encrypted data over a network. An SSL encryption option is provided to disable weak ciphers. Built upon asymmetric encryption technology, SSL is widely accepted for providing authenticated and encrypted communication between clients and servers to prevent eavesdropping across a network.

An SSL-enabled system can perform the following tasks:

- Authenticate itself to an SSL-enabled client
- Allow the two systems to establish an encrypted connection
- **NOTE:** If SSL encryption is set to 256-bit or higher and 168–bit or higher, the cryptography settings for your virtual machine environment (JVM, IcedTea) may require installing the Unlimited Strength Java Cryptography Extension Policy Files to permit usage of iDRAC plugins such as vConsole with this level of encryption. For information about installing the policy files, see the documentation for Java.

iDRAC Web server has a Dell self-signed unique SSL digital certificate by default. You can replace the default SSL certificate with a certificate signed by a well-known Certificate Authority (CA). A Certificate Authority is a business entity that is recognized in the Information Technology industry for meeting high standards of reliable screening, identification, and other important security criteria. Examples of CAs include Thawte and VeriSign. To initiate the process of obtaining a CA-signed certificate, use either iDRAC Web interface or RACADM interface to generate a Certificate Signing Request (CSR) with your



company's information. Then, submit the generated CSR to a CA such as VeriSign or Thawte. The CA can be a root CA or an intermediate CA. After you receive the CA-signed SSL certificate, upload this to iDRAC.

For each iDRAC to be trusted by the management station, that iDRAC's SSL certificate must be placed in the management station's certificate store. Once the SSL certificate is installed on the management stations, supported browsers can access iDRAC without certificate warnings.

You can also upload a custom signing certificate to sign the SSL certificate, rather than relying on the default signing certificate for this function. By importing one custom signing certificate into all management stations, all the iDRACs using the custom signing certificate are trusted. If a custom signing certificate is uploaded when a custom SSL certificate is already in-use, then the custom SSL certificate is disabled and a one-time auto-generated SSL certificate, signed with the custom signing certificate, is used. You can download the custom signing certificate, iDRAC resets and auto-generates a new self-signed SSL certificate. If a self-signed certificate is regenerated, then the trust must be re-established between that iDRAC and the management workstation. Auto-generated SSL certificates are self-signed and have an expiration date of seven years and one day and a start date of one day in the past (for different time zone settings on management stations and the iDRAC).

The iDRAC Web server SSL certificate supports the asterisk character (*) as part of the left-most component of the Common Name when generating a Certificate Signing Request (CSR). For example, *.qa.com, or *.company.qa.com. This is called a wildcard certificate. If a wildcard CSR is generated outside of iDRAC, you can have a signed single wildcard SSL certificate that you can upload for multiple iDRACs and all the iDRACs are trusted by the supported browsers. While connecting to iDRAC Web interface using a supported browser that supports a wildcard certificate, the iDRAC is trusted by the browser. While launching viewers, the iDRACs are trusted by the viewer clients.

Generating a new certificate signing request

A CSR is a digital request to a Certificate Authority (CA) for a SSL server certificate. SSL server certificates allow clients of the server to trust the identity of the server and to negotiate an encrypted session with the server.

After the CA receives a CSR, they review and verify the information the CSR contains. If the applicant meets the CA's security standards, the CA issues a digitally-signed SSL server certificate that uniquely identifies the applicant's server when it establishes SSL connections with browsers running on management stations.

After the CA approves the CSR and issues the SSL server certificate, it can be uploaded to iDRAC. The information used to generate the CSR, stored on the iDRAC firmware, must match the information contained in the SSL server certificate, that is, the certificate must have been generated using the CSR created by iDRAC.

Generating CSR using web interface

To generate a new CSR:

NOTE: Each new CSR overwrites any previous CSR data stored in the firmware. The information in the CSR must match the information in the SSL server certificate. Else, iDRAC does not accept the certificate.

- In the iDRAC Web interface, go to iDRAC Settings > Services > Web Server > SSL certificate, select Generate Certificate Signing Request (CSR) and click Next. The Generate a New Certificate Signing Request page is displayed.
- **2.** Enter a value for each CSR attribute. For more information, see *iDRAC Online Help*.
- 3. Click Generate.

A new CSR is generated. Save it to the management station.

Generating CSR using RACADM

To generate a CSR using RACADM, use the set command with the objects in the iDRAC.Security group, and then use the sslcsrgen command to generate the CSR.

For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.



Automatic Certificate Enrollment

In iDRAC, Automatic certificate enrollment feature enables you for automatic installation and renewal of certificates used by the web server. When this feature is enabled, the existing web server certificate is replaced by a new certificate.

() NOTE:

- Automatic certificate enrollment is a licensed feature and requires Datacenter license.
- Valid NDES (Network Device Enrollment Service) setup is required for issuing the server certificate.

Following are the automatic certificate enrollment configuration parameters:

- Enable / Disable
- SCEP server URL
- Challenge password

(i) **NOTE:** For more information on these parameters, see *iDRAC Online Help*.

Following are the available status for Automatic certificate enrollment:

- Enrolled Automatic certificate enrollment is enabled. Certificate is monitored and new certificate can be issued on expiry.
- Enrolling Intermediate state after Automatic certificate enrollment is enabled.
- Error Problem encountered with NDES server.
- None Default.

NOTE: When you enable Automatic certificate enrollment, web server is restarted and all existing web sessions are logged out.

Uploading server certificate

After generating a CSR, you can upload the signed SSL server certificate to the iDRAC firmware. iDRAC must be reset to apply the certificate. iDRAC accepts only X509, Base 64 encoded Web server certificates. SHA-2 certificates are also supported.

CAUTION: During reset, iDRAC is not available for a few minutes.

Uploading server certificate using web interface

To upload the SSL server certificate:

- In the iDRAC Web interface, go to iDRAC Settings > Connectivity > SSL > SSL certificate, select Upload Server Certificate and click Next.
 - The **Certificate Upload** page is displayed.
- 2. Under File Path, click Browse and select the certificate on the management station.
- 3. Click Apply.
 - The SSL server certificate is uploaded to iDRAC.
- A pop-up message is displayed asking you to reset iDRAC immediately or at a later time. Click Reset iDRAC or Reset iDRAC Later as required.

iDRAC resets and the new certificate is applied. The iDRAC is not available for a few minutes during the reset.

(i) NOTE: You must reset iDRAC to apply the new certificate. Until iDRAC is reset, the existing certificate is active.

Uploading server certificate using RACADM

To upload the SSL server certificate, use the sslcertupload command. For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

If the CSR is generated outside of iDRAC with a private key available, then to upload the certificate to iDRAC:

- 1. Send the CSR to a well-known root CA. CA signs the CSR and the CSR becomes a valid certificate.
- 2. Upload the private key using the remote racadm sslkeyupload command.
- **3.** Upload the signed certificate to iDRAC using the remote racadm sslcertupload command. The new certificate is uploaded iDRAC. A message is displayed asking you to reset iDRAC.



4. Run the racadm racreset command to reset iDRAC.

iDRAC resets and the new certificate is applied. The iDRAC is not available for a few minutes during the reset.

(i) NOTE: You must reset iDRAC to apply the new certificate. Until iDRAC is reset, the existing certificate is active.

Viewing server certificate

You can view the SSL server certificate that is currently being used in iDRAC.

Viewing server certificate using web interface

In the iDRAC Web interface, go to **iDRAC Settings** > **Services** > **Web Server** > **SSL certificate**. The **SSL** page displays the SSL server certificate that is currently in use at the top of the page.

Viewing server certificate using RACADM

To view the SSL server certificate, use the sslcertview command.

For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Uploading custom signing certificate

You can upload a custom signing certificate to sign the SSL certificate. SHA-2 certificates are also supported.

Uploading custom signing certificate using web interface

To upload the custom signing certificate using iDRAC web interface:

- Go to iDRAC Settings > Connectivity > SSL. The SSL page is displayed.
- 2. Under Custom SSL Certificate Signing Certificate, click Upload Signing Certificate. The Upload Custom SSL Certificate Signing Certificate page is displayed.
- **3.** Click **Choose File** and select the custom SSL certificate signing certificate file. Only Public-Key Cryptography Standards #12 (PKCS #12) compliant certificate is supported.
- 4. If the certificate is password protected, in the PKCS#12 Password field, enter the password.
- 5. Click Apply.
- The certificate is uploaded to iDRAC.
- 6. A pop-up message is displayed asking you to reset iDRAC immediately or at a later time. Click **Reset iDRAC** or **Reset iDRAC** immediately or at a later time. Click **Reset iDRAC** or **Reset iDRAC** is a second sec

After iDRAC resets, the new certificate is applied. The iDRAC is not available for a few minutes during the reset.

(i) NOTE: You must reset iDRAC to apply the new certificate. Until iDRAC is reset, the existing certificate is active.

Uploading custom SSL certificate signing certificate using RACADM

To upload the custom SSL certificate signing certificate using RACADM, use the sslcertupload command, and then use the racreset command to reset iDRAC.

For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Downloading custom SSL certificate signing certificate

You can download the custom signing certificate using iDRAC Web interface or RACADM.



Downloading custom signing certificate

To download the custom signing certificate using iDRAC Web interface:

- 1. Go to **iDRAC Settings** > **Connectivity** > **SSL**. The **SSL** page is displayed.
- 2. Under Custom SSL Certificate Signing Certificate, select Download Custom SSL Certificate Signing Certificate and click Next.

A pop-up message is displayed that allows you to save the custom signing certificate to a location of your choice.

Downloading custom SSL certificate signing certificate using RACADM

To download the custom SSL certificate signing certificate, use the sslcertdownload subcommand. For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Deleting custom SSL certificate signing certificate

You can also delete an existing custom signing certificate using iDRAC Web interface or RACADM.

Deleting custom signing certificate using iDRAC web interface

To delete the custom signing certificate using iDRAC web interface:

- Go to iDRAC Settings > Connectivity > SSL. The SSL page is displayed.
- 2. Under Custom SSL Certificate Signing Certificate, select Delete Custom SSL Certificate Signing Certificate and click Next.
- 3. A pop-up message is displayed asking you to reset iDRAC immediately or at a later time. Click **Reset iDRAC** or **Reset iDRAC Later** as required.

After iDRAC resets, a new self-signed certificate is generated.

Deleting custom SSL certificate signing certificate using RACADM

To delete the custom SSL certificate signing certificate using RACADM, use the sslcertdelete subcommand. Then, use the racreset command to reset iDRAC.

For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Configuring multiple iDRACs using RACADM

You can configure one or more iDRACs with identical properties using RACADM. When you query a specific iDRAC using its group ID and object ID, RACADM creates a configuration file from the retrieved information. Import the file to other iDRACs to identically configure them.

() NOTE:

- The configuration file contains information that is applicable for the particular server. The information is organized under various object groups.
- Some configuration files contain unique iDRAC information, such as the static IP address, that you must modify before you import the file into other iDRACs.

You can also use the System Configuration Profile (SCP) to configure multiple iDRACs using RACADM. SCP file contains the component configuration information. You can use this file to apply the configuration for BIOS, iDRAC, RAID, and NIC by importing the file into a target system. For more information, see *XML Configuration Workflow* white paper available at https://www.dell.com/manuals.

To configure multiple iDRACs using the configuration file:



1. Query the target iDRAC that contains the required configuration using the following command:.

racadm get -f <file_name>.xml -t xml -c iDRAC.Embedded.1

The command requests the iDRAC configuration and generates the configuration file.

NOTE: Redirecting the iDRAC configuration to a file using get -f is only supported with the local and remote RACADM interfaces.

(i) NOTE: The generated configuration file does not contain user passwords.

The get command displays all configuration properties in a group (specified by group name and index) and all configuration properties for a user.

2. Modify the configuration file using a text editor, if required.

NOTE: It is recommended that you edit this file with a simple text editor. The RACADM utility uses an ASCII text parser. Any formatting confuses the parser, which may corrupt the RACADM database.

3. On the target iDRAC, use the following command to modify the settings:

```
racadm set -f <file name>.xml -t xml
```

This loads the information into the other iDRAC. You can use set command to synchronize the user and password database with Server Administrator.

4. Reset the target iDRAC using the command: racadm racreset

Disabling access to modify iDRAC configuration settings on host system

You can disable access to modify the iDRAC configuration settings through Local RACADM or iDRAC Settings utility. However, you can view these configuration settings. To do this:

- 1. In iDRAC Web interface, go to iDRAC Settings > Services > Local Configurations.
- 2. Select one or both of the following:
 - **Disable the iDRAC Local Configuration using iDRAC Settings** Disables access to modify the configuration settings in iDRAC Settings utility.
 - **Disable the iDRAC Local Configuration using RACADM** Disables access to modify the configuration settings in Local RACADM.
- 3. Click Apply.

NOTE: If access is disabled, you cannot use Server Administrator or IPMITool to perform iDRAC configurations. However, you can use IPMI Over LAN.



Delegated Authorization using OAuth 2.0

The Delegated Authorization feature allows a user or console to access iDRAC API using OAuth 2.0 JSON Web Tokens (JWT) that the user or console first obtains from an Authorization Server. Once an OAuth JWT has been retrieved, the user or console may use it to invoke iDRAC API. This circumvents the need for specifying username and password to access the API.

NOTE: This feature is only available for DataCenter license. You need to have Configure iDRAC or Configure Users privilege to use this feature.

iDRAC supports configuration of up to 2 Authorization Servers. The configuration requires a user to specify the following Authorization Server details:

- Name A string to identify the Authorization Server on the iDRAC.
- Metadata URL The OpenID Connect compliant URL as advertised by the server.
- HTTPS certificate The server public key the iDRAC should use to communicate with the server.
- Offline Key The JWK set document for the Authorization Server.
- Offline Issuer The issuer string as used in tokens issued by the Authorization Server.

For Online configuration:

- When configuring an Authorization Server, the iDRAC administrator needs to ensure that the iDRAC has online network access to the Authorization Server.
- If iDRAC cannot access the Authorization Server, the configuration fails and a subsequent attempt to access the iDRAC API fails even though a valid token is presented.

For offline configuration:

• iDRAC does not need to communicate with the Auth server, but instead it is configures with the metadata details that it has downloaded offline. When configured offline, iDRAC has public portion of the signing keys and can validate the token without a network connection to the Auth server.



Viewing iDRAC and managed system information

You can view iDRAC and managed system health and properties, hardware and firmware inventory, sensor health, storage devices, network devices, and view and terminate user sessions. For blade servers, you can also view the Flex Address or Remote-Assigned Address (applicable only for MX platforms).

Topics:

- Viewing managed system health and properties
- Configuring Asset Tracking
- Viewing system inventory
- Viewing sensor information
- Monitoring performance index of CPU, memory, and input output modules
- Idle Server Detection
- GPU (Accelerators) Management
- Checking the system for Fresh Air compliance
- Viewing historical temperature data
- Viewing network interfaces available on host OS
- Viewing network interfaces available on host OS using RACADM
- Viewing FlexAddress mezzanine card fabric connections
- Viewing or terminating iDRAC sessions

Viewing managed system health and properties

When you log in to the iDRAC web interface, the **System Summary** page allows you to view the managed system's health, basic iDRAC information, preview the virtual console, add and view work notes, and quickly launch tasks such as power on or off, power cycle, view logs, update and rollback firmware, switch on or switch off the front panel LED, and reset iDRAC.

To access the **System Summary** page, go to **System** > **Overview** > **Summary**. The **System Summary** page is displayed. For more information, see the *iDRAC Online Help*.

You can also view the basic system summary information using the iDRAC Settings utility. To do this, in iDRAC Settings utility, go to **System Summary**. The **iDRAC Settings System Summary** page is displayed. For more information, see the *iDRAC Settings Utility Online Help*.

Configuring Asset Tracking

The Asset Tracking feature in iDRAC provides you the ability to configure various attributes that are related to your server. This includes information such as acquisition, warranty, service, and so on.

NOTE: Asset Tracking in iDRAC is similar to the Asset Tag feature in OpenManage Server Administrator. However, the attribute information has to be entered separately in both these tools for them to report the relevant Asset data.

To configure Asset Tracking:

- 1. In the iDRAC interface, go to **Configuration > Asset Tracking**.
- 2. Click Add Custom Assets to add any additional attributes which are not specified by default on this page.
- 3. Enter all the relevant information of your server asset and click Apply.
- 4. To view the Asset Tracking Report, go to System > Details > Asset Tracking.



Viewing system inventory

You can view information about the hardware and firmware components installed on the managed system. To do this, in iDRAC web interface, go to **System** > **Inventories**. For information about the displayed properties, see the *iDRAC Online Help*.

The Hardware Inventory section displays the information for the following components available on the managed system:

- iDRAC
- RAID controller
- Batteries
- CPUs
- DIMMs
- HDDs
- Backplanes
- Network Interface Cards (integrated and embedded)
- Video card
- SD card
- Power Supply Units (PSUs)
- Fans
- Fibre Channel HBAs
- USB
- NVMe PCIe SSD devices

The Firmware Inventory section displays the firmware version for the following components:

- BIOS
- Lifecycle Controller
- iDRAC
- OS driver pack
- 32-bit diagnostics
- System CPLD
- PERC controllers
- Batteries
- Physical disks
- Power supply
- NIC
- Fibre Channel
- Backplane
- Enclosure
- PCle SSDs

() NOTE:

- Software inventory displays only the last 4 bytes of the firmware version and the Release date information. For example, if the firmware version is FLVDL06, the firmware inventory displays DL06.
- When collecting software inventory using Redfish interface, the Release date information is displayed only for components which support rollback.

NOTE: On the Dell PowerEdge FX2/FX2s servers, the naming convention of the CMC version displayed in the iDRAC GUI differs from that on the CMC GUI. However, the version remains the same.

When you replace any hardware component or update the firmware versions, make sure to enable and run the **Collect System Inventory on Reboot** (CSIOR) option to collect the system inventory on reboot. After a few minutes, log in to iDRAC, and navigate to the **System Inventory** page to view the details. It may take up to 5 minutes for the information to be available depending on the hardware installed on the server.

(i) NOTE: CSIOR option is enabled by default.

NOTE: Configuration changes and firmware updates that are made within the operating system may not reflect properly in the inventory until you perform a server restart.

Click **Export** to export the hardware inventory in an XML format and save it to a location of your choice.



Viewing sensor information

The following sensors help to monitor the health of the managed system:

• **Batteries** — Provides information about the batteries on the system board CMOS and storage RAID On Motherboard (ROMB).

(i) NOTE: The Storage ROMB battery settings are available only if the system has a ROMB with a battery.

- Fan (available only for rack and tower servers) Provides information about the system fans fan redundancy and fans list that display fan speed and threshold values.
- **CPU** Indicates the health and state of the CPUs in the managed system. It also reports processor automatic throttling and predictive failure.
- **Memory** Indicates the health and state of the Dual In-line Memory Modules (DIMMs) present in the managed system.
- Intrusion Provides information about the chassis.
- **Power Supplies** (available only for rack and tower servers) Provides information about the power supplies and the power supply redundancy status.

(i) NOTE: If there is only one power supply in the system, the power supply redundancy is set to **Disabled**.

- **Removable Flash Media** Provides information about the Internal SD Modules; vFlash and Internal Dual SD Module (IDSDM).
 - When IDSDM redundancy is enabled, the following IDSDM sensor status is displayed IDSDM Redundancy Status, IDSDM SD1, IDSDM SD2. When redundancy is disabled, only IDSDM SD1 is displayed.
 - If IDSDM redundancy is initially disabled when the system is powered on or after an iDRAC reset, the IDSDM SD1 sensor status is displayed only after a card is inserted.
 - If IDSDM redundancy is enabled with two SD cards present in the IDSDM, and the status of one SD card is online while the status of the other card is offline. A system reboot is required to restore redundancy between the two SD cards in the IDSDM. After the redundancy is restored, the status of both the SD cards in the IDSDM is online.
 - During the rebuilding operation to restore redundancy between two SD cards present in the IDSDM, the IDSDM status is not displayed since the IDSDM sensors are powered off.

(i) **NOTE:** If the host system is rebooted during IDSDM rebuild operation, the iDRAC does not display the IDSDM information. To resolve this, rebuild IDSDM again or reset the iDRAC.

- System Event Logs (SEL) for a write-protected or corrupt SD card in the IDSDM module are not repeated until they are cleared by replacing the SD card with a writable or good SD card, respectively.
- () NOTE: When iDRAC firmware is updated from versions prior to 3.30.30.30, the iDRAC need to be reset to defaults for IDSDM settings to appear in the Server Administrator's Platform Event Filter.
- **Temperature** Provides information about the system board inlet temperature and exhaust temperature (only applies to rack servers). The temperature probe indicates whether the status of the probe is within the preset warning and critical threshold value.
- Voltage Indicates the status and reading of the voltage sensors on various system components.

The following table provides information about viewing the sensor information using iDRAC web interface and RACADM. For information about the properties that are displayed on the web interface, see the *iDRAC Online Help*.

(i) NOTE: The Hardware Overview page displays data only for sensors present on your system.

Table 17. Sensor information using web interface and RACADM

View sensor information For	Using web interface	Using RACADM
Batteries	Dashboard > System Health > Batteries	Use the getsensorinfo command. For power supplies, you can also use the System.Power.Supply command with the get subcommand. For more information, see the <i>iDRAC</i> <i>RACADM CLI Guide</i> available at https://
Fan	Dashboard > > System Health > Fans	
CPU	Dashboard > System Health > CPU	



Table 17. Sensor information using web interface and RACADM (continued)

View sensor information For	Using web interface	Using RACADM
Memory	Dashboard > System Health > Memory	
Intrusion	Dashboard > System Health > Intrusion	
Power Supplies	> Hardware > Power Supplies	
Removable Flash Media	Dashboard > System Health > Removable Flash Media	
Temperature	Dashboard > System Health > Power/Thermal > Temperatures	
Voltage	Dashboard > System Health > Power/Thermal > Voltages	

Monitoring performance index of CPU, memory, and input output modules

In Dell's 14th generation Dell PowerEdge servers, Intel ME supports Compute Usage Per Second (CUPS) functionality. The CUPS functionality provides real-time monitoring of CPU, memory, and I/O utilization and system-level utilization index for the system. Intel ME allows out-of-band (OOB) performance monitoring and does not consume CPU resources. The Intel ME has a system CUPS sensor that provides computation, memory, and I/O resource utilization values as a CUPS Index. iDRAC monitors this CUPS index for the overall system utilization and also monitors the instantaneous utilization index of the CPU, Memory, and I/O.

() NOTE: CUPS functionality is not supported on following servers:

- PowerEdge R240
- PowerEdge R240xd
- PowerEdge R340
- PowerEdge R6415
- PowerEdge R7415
- PowerEdge R7425
- PowerEdge T140

The CPU and chipset have dedicated Resource monitoring Counters (RMC). The data from these RMCs is queried to obtain utilization information of system resources. The data from RMCs is aggregated by the node manager to measure the cumulative utilization of each of these system resources that is read from iDRAC using existing intercommunication mechanisms to provide data through out-of-band management interfaces.

The Intel sensor representation of performance parameters and index values is for complete physical system. Therefore, the performance data representation on the interfaces is for the complete physical system, even if the system is virtualized and has multiple virtual hosts.

To display the performance parameters, the supported sensors must be present in the server.

The four system utilization parameters are:

- **CPU Utilization** Data from RMCs for each CPU core is aggregated to provide cumulative utilization of all the cores in the system. This utilization is based on time spent in active and inactive states. A sample of RMC is taken every six seconds.
- Memory Utilization RMCs measure memory traffic occurring at each memory channel or memory controller instance. Data from these RMCs is aggregated to measure the cumulative memory traffic across all the memory channels on the system. This is a measure of memory bandwidth consumption and not amount of memory utilization. iDRAC aggregates it for one minute, so it may or may not match the memory utilization that other OS tools, such as **top** in Linux, show. Memory bandwidth utilization that the iDRAC shows is an indication of whether workload is memory intensive or not.
- I/O Utilization There is one RMC per root port in the PCI Express Root Complex to measure PCI Express traffic emanating from or directed to that root port and the lower segment. Data from these RMCs is aggregated for measuring PCI express traffic for all PCI Express segments emanating from the package. This is measure of I/O bandwidth utilization for the system.

Viewing iDRAC and managed system information 117



• System Level CUPS Index — The CUPS index is calculated by aggregating CPU, Memory, and I/O index considering a predefined load factor of each system resource. The load factor depends on the nature of the workload on the system. CUPS Index represents the measurement of the compute headroom available on the server. If the system has a large CUPS Index, then there is limited headroom to place more workload on that system. As the resource consumption decreases, the system's CUPS index decreases. A low CUPS index indicates that there is a large compute headroom and the server can receive new workloads and the server is in a lower power state to reduce power consumption. Workload monitoring can then be applied throughout the data center to provide a high-level and holistic view of the data center's workload, providing a dynamic data center solution.

() NOTE: The CPU, memory, and I/O utilization indexes are aggregated over one minute. Therefore, if there are any

instantaneous spikes in these indexes, they may be suppressed. They are indication of workload patterns not the amount of resource utilization.

The IPMI, SEL, and SNMP traps are generated if the thresholds of the utilization indexes are reached and the sensor events are enabled. The sensor event flags are disabled by default. It can be enabled using the standard IPMI interface.

The required privileges are:

- Login privilege is required to monitor performance data.
- Configure privilege is required for setting warning thresholds and reset historical peaks.
- Login privilege and Enterprise license are required to read historical statics data.

Monitoring performance index of CPU, memory, and input output modules using web interface

To monitor the performance index of CPU, memory, and I/O modules, in the iDRAC web interface, go to **System** > **Performance**.

- System Performance section Displays the current reading and the warning reading for CPU, Memory and I/O utilization index, and system level CUPS index in a graphical view.
- System Performance Historical Data section:
 - Provides the statistics for CPU, memory, IO utilization, and the system level CUPS index. If the host system is powered off, then the graph displays the power off line below 0 percent.
 - You can reset the peak utilization for a particular sensor. Click **Reset Historical Peak**. You must have Configure privilege to reset the peak value.
- Performance Metrics section:
 - Displays status and present reading
 - Displays or specifies the warning threshold utilization limit. You must have server configure privilege to set the threshold values.

For information about the displayed properties, see the *iDRAC Online Help*.

Monitoring performance index for of CPU, memory, and input output modules using RACADM

Use the **SystemPerfStatistics** sub command to monitor performance index for CPU, memory, and I/O modules. For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Idle Server Detection

iDRAC provides out-of-band performance monitoring index of server components like CPU, memory, and I/O.

The history data of the server level CUPS index is used to monitor whether the server is utilized or running idle for long time. If the server is underutilized below certain threshold for a defined span of interval (in hours), then it will be reported as idle server.

This feature is only supported on Intel platforms with CUPS ability. AMD and Intel platforms without CUPS capability do not support this feature.

() NOTE:

- This feature requires Datacenter license.
- 118 Viewing iDRAC and managed system information



 To read the configurations of Idle Server Configuration parameters, you need Login privilege and to modify the parameters you need iDRAC Configure privilege.

To view or modify the parameters, navigate to **Configuration** > **System Settings**.

Idle server detection is reported based on following parameters:

- Idle Server Threshold (%) This is set to 20% by default and can be configured from 0 to 50%. The reset operation sets the threshold to 20%.
- Idle Server Scan Interval (in hours) This is the time period over which the hourly samples are collected to determine the idle server. This is set to 240 hours by default and can be configured from 1 to 9000 hours. The reset operation sets the interval to 240 hours.
- Server Utilization Percentile (%) The utilization percentile value can be set to 80 to 100%. The default value is 80%. If the 80% of the hourly samples falls below utilization threshold, then it is considered as idle server.

Modifying idle Server Detection parameters using RACADM

racadm get system.idleServerDetection

Modifying idle Server Detection parameters using Redfish

https://<iDRAC IP>/redfish/v1/Managers/System.Embedded.1/Attributes

Modifying idle Server Detection parameters using WSMAN

winrm e http://schemas.dmtf.org/wbem/wscim/1/cim-schema/2/root/dcim/DCIM_SystemAttribute -u:root -p:calvin -r:https://<iDRAC IP>/wsman -SkipCNcheck -SkipCAcheck -encoding:utf-8 -a:basic

iNOTE: iDRAC GUI doesn't support to view or modify the attributes.

GPU (Accelerators) Management

Dell PowerEdge servers are shipped with Graphics Processing Unit (GPU). GPU management enables you to view the various GPUs connected to the system and also monitor power, temperature, and thermal information for the GPUs.

NOTE: This is a licensed feature and is available only with iDRAC Datacenter and Enterprise licenses. Below properties require Datacenter/Enterprise license, other properties are listed even without these license:

- Thermal Metrics:
 - GPU Target Temperature
 - Minimum GPU HW Slowdown Temperature
 - GPU Shutdown Temperature
 - Maximum Memory Operating temperature
 - Maximum GPU Operating Temperature
 - Thermal Alert State
 - Power Brake State
- Power Metrics:
 - Power Supply Status
 - Board Power Supply Status
- Telemetry All GPU telemetry reports data

(i) NOTE: GPU properties will not be listed for Embedded GPU cards and the Status is marked as **Unknown**.



GPU has to be in ready state before the command fetches the data. GPUStatus field in Inventory shows the availability of the GPU and whether GPU device is responding or not. If the GPU status is ready, GPUStatus shows OK, otherwise the status shows Unavailable.

The GPU offers multiple health parameters which can be pulled through the SMBPB interface of the NVIDIA controllers. This feature is limited only to NVIDIA cards. Following are the health parameters retrieved from the GPU device:

- Power
- Temperature
- Thermal

NOTE: This feature is only limited to NVIDIA cards. This information is not available for any other GPU that the server may support. The interval for polling the GPU cards over the PBI is 5 seconds.

The host system must have the NVIDIA driver installed and running for the Power consumption, GPU target temperature, Min GPU slowdown temperature, GPU shutdown temperature, Max memory operating temperature, and Max GPU operating temperature features to be available. These values are shown as **N/A** if the GPU driver is not installed.

In Linux, when the card is unused, the driver down-trains the card and unloads in order to save power. In such cases, the Power consumption, GPU target temperature, Min GPU slowdown temperature, GPU shutdown temperature, Max memory operating temperature, and Max GPU operating temperature features are not available. Persistent mode should be enabled for the device to avoid unload. You can use nvidia-smi tool to enable this using the command nvidia-smi -pm 1.

You can generate GPU reports using Telemetry. For more information on telemetry feature, see Telemetry Streaming

NOTE: In Racadm, You may see dummy GPU entries with empty values. This may happen if device is not ready to respond when iDRAC queries the GPU device for the information. Perform iDRAC racrest operation to resolve this issue.

FPGA Monitoring

Field-programmable Gate Array (FPGA) devices needs real-time temperature sensor monitoring as it generates significant heat when in use. Perform the following steps to get FPGA inventory information:

- Power off the server.
- Install FPGA device on the riser card.
- Power on the server.
- Wait until POST is complete.
- Login to iDRAC GUI.
- Navigate to System > Overview > Accelerators. You can see both GPU and FPGA sections.
- Expand the specific FPGA component to see the following sensor information:
- Power consumption
- Temperature details

(i) NOTE: You must have iDRAC Login privilege to access FPGA information.

NOTE: Power consumption sensors are available only for the supported FPGA cards and is available only with Datacenter license.

Checking the system for Fresh Air compliance

Fresh Air cooling directly uses outside air to cool systems in the data center. Fresh Air compliant systems can operate above its normal ambient operating range (temperatures up to 113 °F (45 °C)).

NOTE: Some servers or certain configurations of a server may not be Fresh Air compliant. See the specific server manual for details related to Fresh Air compliance or contact Dell for more details.

To check the system for Fresh Air compliance:

- In the iDRAC Web interface, go to System > Overview > Cooling > Temperature overview. The Temperature overview page is displayed.
- 2. See the Fresh Air section that indicates whether the server is fresh air compliant or not.



Viewing historical temperature data

You can monitor the percentage of time the system has operated at ambient temperature that is greater than the normally supported fresh air temperature threshold. The system board temperature sensor reading is collected over a period of time to monitor the temperature. The data collection starts when the system is first powered on after it is shipped from the factory. The data is collected and displayed for the duration when the system is powered on. You can track and store the monitored temperature for the last seven years.

NOTE: You can track the temperature history even for systems that are not Fresh-Air compliant. However, the threshold limits and fresh air related warnings generated are based on fresh air supported limits. The limits are 42°C for warning and 47°C for critical. These values correspond to 40°C and 45°C fresh air limits with 2°C margin for accuracy.

Two fixed temperature bands are tracked that are associated to fresh air limits:

- Warning band Consists of the duration a system has operated above the temperature sensor warning threshold (42°C). The system can operate in the warning band for 10% of the time for 12 months.
- Critical band Consists of the duration a system has operated above the temperature sensor critical threshold (47°C). The system can operate in the critical band for 1% of the time for 12 months which also increments time in the warning band.

The collected data is represented in a graphical format to track the 10% and 1% levels. The logged temperature data can be cleared only before shipping from the factory.

An event is generated if the system continues to operate above the normally supported temperature threshold for a specified operational time. If the average temperature over the specified operational time is greater than or equal to the warning level (> = 8%) or the critical level (> = 0.8%), an event is logged in the Lifecycle Log and the corresponding SNMP trap is generated. The events are:

- Warning event when the temperature was greater than the warning threshold for duration of 8% or more in the last 12 months.
- Critical event when the temperature was greater than the warning threshold for duration of 10% or more in the last 12 months.
- Warning event when the temperature was greater than the critical threshold for duration of 0.8% or more in the last 12 months.
- Critical event when the temperature was greater than the critical threshold for duration of 1% or more in the last 12 months.

You can also configure iDRAC to generate additional events. For more information, see the Setting alert recurrence event section.

Viewing historical temperature data using iDRAC web interface

To view historical temperature data:

- In the iDRAC Web interface, go to System > Overview > Cooling > Temperature overview. The Temperature overview page is displayed.
- 2. See the **System Board Temperature Historical Data** section that provides a graphical display of the stored temperature (average and peak values) for the last day, last 30 days, and last year.

For more information, see the *iDRAC Online Help*.

(i) NOTE: After an iDRAC firmware update or iDRAC reset, some temperature data may not be displayed in the graph.

NOTE: WX3200 AMD GPU card currently doesnot support I2C interface for temperature sensors. Hence, temperature readings will not be available for this card from iDRAC interfaces.

Viewing historical temperature data using RACADM

To view historical data using RACADM, use the inlettemphistory command.

For more information, see the iDRAC RACADM CLI Guide available at https://www.dell.com/idracmanuals.



Configuring warning threshold for inlet temperature

You can modify the minimum and maximum warning threshold values for the system board inlet temperature sensor. If reset to default action is performed, the temperature thresholds are set to the default values. You must have Configure user privilege to set the warning threshold values for the inlet temperature sensor.

Configuring warning threshold for inlet temperature using web interface

To configure warning threshold for inlet temperature:

- 1. In the iDRAC Web interface, go to **System** > **Overview** > **Cooling** > **Temperature overview**. The **Temperature overview** page is displayed.
- In the Temperature Probes section, for the System Board Inlet Temp, enter the minimum and maximum values for the Warning Threshold in Centigrade or Fahrenheit. If you enter the value in centigrade, the system automatically calculates and displays the Fahrenheit value. Similarly, if you enter Fahrenheit, the value for Centigrade is displayed.
- 3. Click Apply.

The values are configured.

NOTE: Changes to default thresholds are not reflected in the historical data chart since the chart limits are for fresh air limit values only. Warnings for exceeding the custom thresholds are different from warning associated to exceeding fresh air thresholds.

Viewing network interfaces available on host OS

You can view information about all the network interfaces that are available on the host operating system such as the IP addresses that are assigned to the server. The iDRAC Service Module provides this information to iDRAC. The OS IP address information includes the IPv4 and IPv6 addresses, MAC address, Subnet mask or prefix length, the FQDD of the network device, network interface name, network interface description, network interface status, network interface type (Ethernet, tunnel, loopback, and so on.), Gateway address, DNS server address, and DHCP server address.

(i) NOTE: This feature is available with iDRAC Express and iDRAC Enterprise/Datacenter licenses.

To view the OS information, make sure that:

- You have Login privilege.
- iDRAC Service Module is installed and running on the host operating system.
- OS Information option is enabled in the **iDRAC Settings** > **Overview** > **iDRAC Service Module** page.

iDRAC can display the IPv4 and IPv6 addresses for all the interfaces configured on the Host OS.

Depending on how the Host OS detects the DHCP server, the corresponding IPv4 or IPv6 DHCP server address may not be displayed.

Viewing network interfaces available on host OS using web interface

To view the network interfaces available on the host OS using Web interface:

- 1. Go to System > Host OS > Network Interfaces.
- The Network Interfaces page displays all the network interfaces that are available on the host operating system.
- To view the list of network interfaces associated with a network device, from the Network Device FQDD drop-down menu, select a network device and click Apply. The OS IP details are displayed in the Host OS Network Interfaces section.
- The US IP details are displayed in the **Host US Network Interfaces** sect
- 3. From the Device FQDD column, click on the network device link. The corresponding device page is displayed from the Hardware > Network Devices section, where you can view the device details. For information about the properties, see the *iDRAC Online Help*.
- 4. Click the 🛨 icon to display more details.

Similarly, you can view the host OS network interface information associated with a network device from the **Hardware** > **Network Devices** page. Click **View Host OS Network Interfaces**.

122 Viewing iDRAC and managed system information



NOTE: For the ESXi host OS in the iDRAC Service Module v2.3.0 or later, the **Description** column in the **Additional Details** list is displayed in the following format:

<List-of-Uplinks-Configured-on-the-vSwitch>/<Port-Group>/<Interface-name>

Viewing network interfaces available on host OS using RACADM

Use the gethostnetworkinterfaces command to view the network interfaces available on the host operating systems using RACADM. For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Viewing FlexAddress mezzanine card fabric connections

In blade servers, FlexAddress allows the use of persistent, chassis-assigned World Wide Names and MAC addresses (WWN/MAC) for each managed server port connection.

You can view the following information for each installed embedded Ethernet and optional mezzanine card port:

- Fabrics to which the cards are connected.
- Type of fabric.
- Server-assigned, chassis-assigned, or remotely assigned MAC addresses.

To view the Flex Address information in iDRAC, configure and enable the Flex Address feature in Chassis Management Controller (CMC). For more information, see the *Chassis Management Controller User's Guide* available at https://www.dell.com/cmcmanuals. Any existing Virtual Console or Virtual Media session terminates if the FlexAddress setting is enabled or disabled.

NOTE: To avoid errors that may lead to an inability to turn on the managed system, you *must* have the correct type of mezzanine card installed for each port and fabric connection.

The FlexAddress feature replaces the server–assigned MAC addresses with chassis–assigned MAC addresses and is implemented for iDRAC along with blade LOMs, mezzanine cards and I/O modules. The iDRAC FlexAddress feature supports preservation of slot specific MAC address for iDRACs in a chassis. The chassis–assigned MAC address is stored in CMC non–volatile memory and is sent to iDRAC during an iDRAC boot or when CMC FlexAddress is enabled.

If CMC enables chassis-assigned MAC addresses, iDRAC displays the MAC address on any of the following pages:

- System > Details > iDRAC Details.
- System > Server > WWN/MAC.
- iDRAC Settings > Overview > Current Network Settings.

CAUTION: With FlexAddress enabled, if you switch from a server-assigned MAC address to a chassis-assigned MAC address and vice-versa, iDRAC IP address also changes.

Viewing or terminating iDRAC sessions

You can view the number of users currently logged in to iDRAC and terminate the user sessions.

Terminating iDRAC sessions using web interface

The users who do not have administrative privileges must have Configure iDRAC privilege to terminate iDRAC sessions using iDRAC Web interface.

To view and terminate the iDRAC sessions:

1. In the iDRAC Web interface, go to **iDRAC Settings** > **users** > **Sessions**.



The **Sessions** page displays the session ID, username, IP address, and session type. For more information about these properties, see the *iDRAC Online Help*.

2. To terminate the session, under the **Terminate** column, click the Trashcan icon for a session.

Terminating iDRAC sessions using RACADM

You must have administrator privileges to terminate iDRAC sessions using RACADM.

To view the current user sessions, use the getssninfo command.

To terminate a user session, use the closessn command.

For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

124 Viewing iDRAC and managed system information



Setting up iDRAC communication

You can communicate with iDRAC using any of the following modes:

- iDRAC Web Interface
- Serial connection using DB9 cable (RAC serial or IPMI serial) For rack and tower servers only
- IPMI Serial Over LAN
- IPMI Over LAN
- Remote RACADM
- Local RACADM
- Remote Services
- **NOTE:** To ensure that Local RACADM import or export commands work properly, ensure that the USB mass-storage host is enabled in the operating system. For information about enabling USB storage host, see the documentation for your operating system.

The following table provides an overview of the supported protocols, supported commands, and pre-requisites:

Table 18. Communication modes — summary

Mode of Communication	Supported Protocol	Supported Commands	Pre-requisite
iDRAC Web Interface	Internet Protocol (https)	N/A	Web Server
Serial using Null modem	Serial Protocol	RACADM	Part of iDRAC firmware
		IPMI	RAC Serial or IPMI Serial is enabled
IPMI Serial Over LAN	Intelligent Platform Management Bus protocol SSH	IPMI	IPMITool is installed and IPMI Serial Over LAN is enabled
IPMI over LAN	Intelligent Platform Management Bus protocol	IPMI	IPMITool is installed and IPMI Settings is enabled
Remote RACADM	https	Remote RACADM	Remote RACADM is installed and enabled
Firmware RACADM	SSH	Firmware RACADM	Firmware RACADM is installed and enabled
Local RACADM	IPMI	Local RACADM	Local RACADM is installed
Remote Services ¹	WSMan	WinRM (Windows)	WinRM is installed (Windows)
		OpenWSMan (Linux)	(Linux)
	Redfish	Various browser plug-ins, CURL (Windows and Linux), Python request and JSON modules	Plug-ins, CURL, Python modules are installed
[1] For more information, see the Lifecycle Controller User's Guide available at https://www.dell.com/idracmanuals .			

Topics:

- Communicating with iDRAC through serial connection using DB9 cable
- Switching between RAC serial and serial console while using DB9 cable
- Communicating with iDRAC using IPMI SOL
- Communicating with iDRAC using IPMI over LAN



- Enabling or disabling remote RACADM
- Disabling local RACADM
- Enabling IPMI on managed system
- Configuring Linux for serial console during boot in RHEL 6
- Configuring serial terminal in RHEL 7
- Supported SSH cryptography schemes

Communicating with iDRAC through serial connection using DB9 cable

You can use any of the following communication methods to perform systems management tasks through serial connection to rack and tower servers:

- RAC Serial
- IPMI Serial Direct Connect Basic mode and Direct Connect Terminal mode
- **NOTE:** In case of blade servers, the serial connection is established through the chassis. For more information, see the *Chassis Management Controller User's Guide* available at https://www.dell.com/cmcmanuals (not applicable for MX platforms) *OME Modular for PowerEdge MX7000 Chassis User's Guide* available at https://www.dell.com/openmanagemanuals (applicable for MX platforms).

To establish the serial connection:

- **1.** Configure the BIOS to enable serial connection.
- 2. Connect the Null Modem DB9 cable from the management station's serial port to the managed system's external serial connector.

(i) NOTE: Server power cycle is required from vConsole or GUI for any change in Baud-rate.

NOTE: If iDRAC serial connection authentication is disabled, then iDRAC racreset is required for any change in BAUD-rate.

- **3.** Make sure that the management station's terminal emulation software is configured for serial connection using any of the following:
 - Linux Minicom in an Xterm
 - Hilgraeve's HyperTerminal Private Edition (version 6.3)

Based on where the managed system is in its boot process, you can see either the POST screen or the operating system screen. This is based on the configuration: SAC for Windows and Linux text mode screens for Linux.

4. Enable RAC serial or IPMI serial connections in iDRAC.

Configuring BIOS for serial connection

To configure BIOS for Serial Connection:

(i) NOTE: This is applicable only for iDRAC on rack and tower servers.

- 1. Turn on or restart the system.
- 2. Press F2.
- **3.** Go to System BIOS Settings > Serial Communication.
- 4. Select External Serial Connector to Remote Access device.
- 5. Click Back, click Finish, and then click Yes.
- 6. Press Esc to exit System Setup.

Enabling RAC serial connection

After configuring serial connection in BIOS, enable RAC serial in iDRAC.

126 Setting up iDRAC communication



(i) NOTE: This is applicable only for iDRAC on rack and tower servers.

Enabling RAC serial connection using web interface

To enable RAC serial connection:

- In the iDRAC Web interface, go to iDRAC Settings > Network > Serial. The Serial page is displayed.
- 2. Under RAC Serial, select Enabled and specify the values for the attributes.
- 3. Click Apply.
 - The RAC serial settings are configured.

Enabling RAC serial connection using RACADM

To enable RAC serial connection using RACADM, use the set command with the object in the iDRAC. Serial group.

Enabling IPMI serial connection basic and terminal modes

To enable IPMI serial routing of BIOS to iDRAC, configure IPMI Serial in any of the following modes in iDRAC:

(i) NOTE: This is applicable only for iDRAC on rack and tower servers.

 IPMI basic mode — Supports a binary interface for program access, such as the IPMI shell (ipmish) that is included with the Baseboard Management Utility (BMU). For example, to print the System Event Log using ipmish via IPMI Basic mode, run the following command:

ipmish -com 1 -baud 57600 -flow cts -u <username> -p <password> sel get

(i) NOTE: The default iDRAC user name and password are provided on the system badge.

- IPMI terminal mode Supports ASCII commands that are sent from a serial terminal. This mode supports limited number of commands (including power control) and raw IPMI commands that are typed as hexadecimal ASCII characters. It allows you to view the operating system boot sequences up to BIOS, when you login to iDRAC through SSH. You need to logout from the IPMI terminal using [sys pwd -x], below are the example for IPMI Terminal mode commands.
 - o [sys tmode]

```
o [sys pwd -u root calvin]
```

- o [sys health query -v]
- [18 00 01]
- o [sys pwd -x]

Enabling serial connection using web interface

Make sure to disable the RAC serial interface to enable IPMI Serial.

To configure IPMI Serial settings:

- 1. In the iDRAC Web interface, go to iDRAC Settings > Connectivity > Serial.
- 2. Under IPMI Serial, specify the values for the attributes. For information about the options, see the iDRAC Online Help.
- 3. Click Apply.

Enabling serial connection IPMI mode using RACADM

To configure the IPMI mode, disable the RAC serial interface and then enable the IPMI mode.

```
racadm set iDRAC.Serial.Enable 0
racadm set iDRAC.IPMISerial.ConnectionMode <n>
```

n=0 — Terminal Mode



n=1 — Basic Mode

Enabling serial connection IPMI serial settings using RACADM

1. Change the IPMI serial-connection mode to the appropriate setting using the command.

racadm set iDRAC.Serial.Enable 0

2. Set the IPMI Serial baud rate using the command.

```
racadm set iDRAC.IPMISerial.BaudRate <baud rate>
```

Parameter	Allowed values (in bps)
<baud_rate></baud_rate>	9600, 19200, 57600, and 115200.

3. Enable the IPMI serial hardware flow control using the command.

racadm set iDRAC.IPMISerial.FlowContro 1

4. Set the IPMI serial channel minimum privilege level using the command.

racadm set iDRAC.IPMISerial.ChanPrivLimit <level>

Parameter	Privilege level
<level> = 2</level>	User
<level> = 3</level>	Operator
<level> = 4</level>	Administrator

5. Ensure that the serial MUX (external serial connector) is set correctly to the remote access device in the BIOS Setup program to configure BIOS for serial connection.

For more information about these properties, see the IPMI 2.0 specification.

Additional settings for ipmi serial terminal mode

This section provides additional configuration settings for IPMI serial terminal mode.

Configuring additional settings for IPMI serial terminal mode using web interface

To set the Terminal Mode settings:

- In the iDRAC Web interface, go to iDRAC Settings > Connectivity > Serial. The Serial page is displayed.
- 2. Enable IPMI serial.
- **3.** Click **Terminal Mode Settings.** The **Terminal Mode Settings** page is displayed.
- **4.** Specify the following values:
 - Line editing
 - Delete control
 - Echo Control
 - Handshaking control
 - New line sequence
 - Input new line sequences

For information about the options, see the *iDRAC Online Help*.


5. Click Apply.

The terminal mode settings are configured.

6. Make sure that the serial MUX (external serial connector) is set correctly to the remote access device in the BIOS Setup program to configure BIOS for serial connection.

Configuring additional settings for IPMI serial terminal mode using RACADM

To configure the Terminal Mode settings, use the set command with the objects in the idrac.ipmiserial group. For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Switching between RAC serial and serial console while using DB9 cable

iDRAC supports Escape key sequences that allow switching between RAC Serial Interface communication and Serial Console on rack and tower servers.

Switching from serial console to RAC serial

To switch to RAC Serial Interface communication mode when in Serial Console Mode, press Esc+Shift, 9.

The key sequence directs you to the iDRAC Login prompt (if the iDRAC is set to RAC Serial mode) or to the Serial Connection mode where terminal commands can be issued if iDRAC is set to IPMI Serial Direct Connect Terminal Mode.

Switching from RAC serial to serial console

To switch to Serial Console Mode when in RAC Serial Interface Communication Mode, press Esc+Shift, Q.

When in terminal mode, to switch the connection to the Serial Console mode, press Esc+Shift, Q.

To go back to the terminal mode use, when connected in Serial Console mode, press Esc+Shift, 9.

Communicating with iDRAC using IPMI SOL

IPMI Serial Over LAN (SOL) allows a managed system's text-based console serial data to be redirected over iDRAC's dedicated or shared out-of-band ethernet management network. Using SOL you can:

- Remotely access operating systems with no time-out.
- Diagnose host systems on Emergency Management Services (EMS) or Special Administrator Console (SAC) for Windows or Linux shell.
- View the progress of a servers during POST and reconfigure the BIOS setup program.

To setup the SOL communication mode:

- 1. Configure BIOS for serial connection.
- **2.** Configure iDRAC to Use SOL.
- 3. Enable a supported protocol (SSH, IPMItool).

Configuring BIOS for serial connection

(i) NOTE: This is applicable only for iDRAC on rack and tower servers.

- 1. Turn on or restart the system.
- 2. Press F2.
- 3. Go to System BIOS Settings > Serial Communication.
- **4.** Specify the following values:



- Serial Communication On With Console Redirection
- Serial Port Address COM2.
 NOTE: You can set the serial communication field to On with serial redirection via com1 if serial device2 in the serial port address field is also set to com1.
- External serial connector Serial device 2
- Failsafe Baud Rate 115200
- Remote Terminal Type VT100/VT220
- Redirection After Boot Enabled
- 5. Click Back and then click Finish.
- 6. Click Yes to save the changes.
- 7. Press <Esc> to exit System Setup.
 - **NOTE:** BIOS sends the screen serial data in 25 x 80 format. The SSH window that is used to invoke the console com2 command must be set to 25 x 80. Then, the redirected screen appears correctly.
 - () NOTE: If the boot loader or operating system provides serial redirection such as GRUB or Linux, then the BIOS **Redirection After Boot** setting must be disabled. This is to avoid potential race condition of multiple components accessing the serial port.

Configuring iDRAC to use SOL

You can specify the SOL settings in iDRAC using Web interface, RACADM, or iDRAC Settings utility.

Configuring iDRAC to use SOL using iDRAC web interface

To configure IPMI Serial over LAN (SOL):

- In the iDRAC Web interface, go to iDRAC Settings > Connectivity > Serial Over LAN. The Serial over LAN page is displayed.
- **2.** Enable SOL, specify the values, and click **Apply**. The IPMI SOL settings are configured.
- **3.** To set the character accumulate interval and the character send threshold, select **Advanced Settings**. The **Serial Over LAN Advanced Settings** page is displayed.
- 4. Specify the values for the attributes and click Apply.

The IPMI SOL advanced settings are configured. These values help to improve the performance.

For information about the options, see the *iDRAC Online Help*.

Configuring iDRAC to use SOL using RACADM

To configure IPMI Serial over LAN (SOL):

1. Enable IPMI Serial over LAN using the command.

racadm set iDRAC.IPMISol.Enable 1

2. Update the IPMI SOL minimum privilege level using the command.

racadm set iDRAC.IPMISol.MinPrivilege <level>

Parameter	Privilege level
<level> = 2</level>	User
<level> = 3</level>	Operator
<level> = 4</level>	Administrator



NOTE: To activate IPMI SOL, you must have the minimum privilege defined in IMPI SOL. For more information, see the IPMI 2.0 specification.

3. Update the IPMI SOL baud rate using the command.

racadm set iDRAC.IPMISol.BaudRate <baud_rate>

NOTE: To redirect the serial console over LAN, make sure that the SOL baud rate is identical to the managed system's baud rate.

Parameter	Allowed values (in bps)
<baud_rate></baud_rate>	9600, 19200, 57600, and 115200.

4. Enable SOL for each user using the command.

racadm set iDRAC.Users.<id>.SolEnable 2

Parameter	Description	
<id></id>	Unique ID of the user	

NOTE: To redirect the serial console over LAN, ensure that the SOL baud rate is identical to the baud rate of the managed system.

Enabling supported protocol

The supported protocols are IPMI and SSH.

Enabling supported protocol using web interface

To enable SSH, go to **iDRAC Settings** > Services and select Enabled for SSH.

To enable IPMI, go to **iDRAC Settings** > **Connectivity** and select **IPMI Settings**. Make sure that the **Encryption Key** value is all zeroes or press the backspace key to clear and change the value to NULL characters.

Enabling supported protocol using RACADM

To enable the SSH, use the following command.

SSH

```
racadm set iDRAC.SSH.Enable 1
```

To change the SSH port

racadm set iDRAC.SSH.Port <port number>

You can use tools such as:

- IPMItool for using IPMI protocol
- Putty/OpenSSH for using SSH protocol

SOL using IPMI protocol

The IPMI-based SOL utility and IPMItool use RMCP+ delivered using UDP datagrams to port 623. The RMCP+ provides improved authentication, data integrity checks, encryption, and the ability to carry multiple types of payloads while using IPMI 2.0. For more information, see http://ipmitool.sourceforge.net/manpage.html.

Setting up iDRAC communication 131



The RMCP+ uses a 40-character hexadecimal string (characters 0-9, a-f, and A-F) encryption key for authentication. The default value is a string of 40 zeros.

An RMCP+ connection to iDRAC must be encrypted using the encryption key (Key Generator Key). You can configure the encryption key using the iDRAC web interface or iDRAC Settings utility.

To start SOL session using IPMItool from a management station:

(i) NOTE: If required, you can change the default SOL time-out at iDRAC Settings > Services.

- 1. Install IPMITool from the *Dell Systems Management Tools and Documentation* DVD. For installation instructions, see the *Software Quick Installation Guide*.
- 2. At the command prompt (Windows or Linux), run the following command to start SOL from iDRAC:

```
ipmitool -H <iDRAC-ip-address> -I lanplus -U <login name> -P <login password> sol activate
```

This command connected the management station to the managed system's serial port.

3. To quit a SOL session from IPMItool, press \sim and then . (period).

(i) NOTE: If a SOL session does not terminate, reset iDRAC and allow up to two minutes to complete booting.

- **NOTE:** IPMI SOL session may terminate while copying large input text from a client running Windows OS to a host running Linux OS. To avoid the session from getting terminated abruptly, convert any large text to a UNIX-based line ending.
- **NOTE:** If a SOL session created using RACADM tool exists, starting another SOL session using IPMI tool will not show any notification or error about the existing sessions.
- **NOTE:** Due to windows OS settings, SOL session connected through ssh and IPMI tool may go to blank screen after booting. Disconnect and Re-connect the SOL Session again to get back SAC prompt.

SOL using SSH

Secure Shell (SSH) is network protocol used to perform command line communications to iDRAC. You can parse remote RACADM commands through this interface.

SSH has improved security. iDRAC only supports SSH version 2 with password authentication, and is enabled by default. iDRAC supports up to two to four SSH sessions at a time.

- () NOTE: Starting iDRAC version 4.40.00.00, telnet feature is removed from iDRAC, so any related attribute registry properties are obsoleted. While some of these properties are still available in iDRAC to keep backward compatibility with existing console applications and scripts, the corresponding settings are ignored by iDRAC firmware.
- (i) NOTE: While establishing SSH connection, a security message is displayed- 'Further Authentication required'. even though 2FA is disabled.
- **NOTE:** For MX platforms, one SSH session will be used for iDRAC communication. If all the sessions are in use, then iDRAC will not launch until one session is free.

Use open-source programs such as PuTTY or OpenSSH that support SSH on a management station to connect to iDRAC.

NOTE: Run OpenSSH from a VT100 or ANSI terminal emulator on Windows. Running OpenSSH at the Windows command prompt does not result in full functionality (that is, some keys do not respond and no graphics are displayed).

Before using SSH to communicate with iDRAC, make sure to:

- 1. Configure BIOS to enable Serial Console.
- 2. Configure SOL in iDRAC.
- 3. Enable SSH using iDRAC Web interface or RACADM.
 - SSH (port 22) client <--> WAN connection <--> iDRAC

The IPMI-based SOL that uses SSH protocol eliminates the need for an additional utility because the serial to network translation happens within iDRAC. The SSH console that you use must be able to interpret and respond to the data arriving



from the serial port of the managed system. The serial port usually attaches to a shell that emulates an ANSI- or VT100/VT220-terminal. The serial console is automatically redirected to the SSH.

Using SOL from PuTTY on Windows

(i) NOTE: If required, you can change the default SSH time-out at **iDRAC Settings** > Services.

To start IPMI SOL from PuTTY on a Windows management station:

1. Run the following command to connect to iDRAC

putty.exe [-ssh] <login name>@<iDRAC-ip-address> <port number>

(i) NOTE: The port number is optional. It is required only when the port number is reassigned.

2. Run the command console com2 or connect to start SOL and boot the managed system.

A SOL session from the management station to the managed system using the SSH protocol is opened. To access the iDRAC command-line console, follow the ESC key sequence. Putty and SOL connection behavior:

- While accessing the managed system through putty during POST, if the Function keys and keypad option on putty is set to:
 - VT100+ F2 passes, but F12 cannot pass.
 - ESC[n~ F12 passes, but F2 cannot pass.
- In Windows, if the Emergency Management System (EMS) console is opened immediately after a host reboot, the Special Admin Console (SAC) terminal may get corrupted. Quit the SOL session, close the terminal, open another terminal, and start the SOL session using the same command.
- **NOTE:** Due to windows OS settings, SOL session connected through ssh and IPMI tool may go to blank screen after booting. Disconnect and Re-connect the SOL Session again to get back SAC prompt.

Using SOL from OpenSSH on Linux

To start SOL from OpenSSH on a Linux management station:

(i) NOTE: If required, you can change the default SSH session time-out at iDRAC Settings > Services.

- 1. Start a shell.
- 2. Connect to iDRAC using the following command: ssh <iDRAC-ip-address> -I <login name>
- **3.** Enter one of the following commands at the command prompt to start SOL:
 - connect
 - console com2

This connects iDRAC to the managed system's SOL port. Once a SOL session is established, iDRAC command line console is not available. Follow the escape sequence correctly to open the iDRAC command line console. The escape sequence is also printed on the screen as soon as a SOL session is connected. When the managed system is off, it takes sometime to establish the SOL session.

(i) NOTE: You can use console com1 or console com2 to start SOL. Reboot the server to establish the connection.

The console -h com2 command displays the contents of the serial history buffer before waiting for input from the keyboard or new characters from the serial port.

The default (and maximum) size of the history buffer is 8192 characters. You can set this number to a smaller value using the command:

racadm set iDRAC.Serial.HistorySize <number>

4. Quit the SOL session to close an active SOL session.



Disconnecting SOL session in iDRAC command line console

The commands to disconnect a SOL session are based on the utility. You can exit the utility only when a SOL session is completely terminated.

To disconnect a SOL session, terminate the SOL session from the iDRAC command line console.

 To quit SOL redirection, press Enter, Esc, T. The SOL session closes.

If a SOL session is not terminated completely in the utility, other SOL sessions may not be available. To resolve this, terminate the command line console in the Web interface under **iDRAC Settings** > **Connectivity** > **Serial Over LAN**.

Communicating with iDRAC using IPMI over LAN

You must configure IPMI over LAN for iDRAC to enable or disable IPMI commands over LAN channels to any external systems. If IPMI over LAN is not configured, then external systems cannot communicate with the iDRAC server using IPMI commands.

(i) NOTE: IPMI also supports IPv6 address protocol for Linux-based operating systems.

Configuring IPMI over LAN using web interface

To configure IPMI over LAN:

- In the iDRAC Web interface, go to iDRAC Settings > Connectivity. The Network page is displayed.
- 2. Under IPMI Settings, specify the values for the attributes and click Apply.

For information about the options, see the *iDRAC Online Help*.

The IPMI over LAN settings are configured.

Configuring IPMI over LAN using iDRAC settings utility

To configure IPMI over LAN:

- In the iDRAC Settings Utility, go to Network. The iDRAC Settings Network page is displayed.
- 2. For IPMI Settings, specify the values.

For information about the options, see the *iDRAC Settings Utility Online Help*.

3. Click **Back**, click **Finish**, and then click **Yes**. The IPMI over LAN settings are configured.

Configuring IPMI over LAN using RACADM

1. Enable IPMI over LAN.

```
racadm set iDRAC.IPMILan.Enable 1
```

- (i) **NOTE:** This setting determines the IPMI commands that are executed using IPMI over LAN interface. For more information, see the IPMI 2.0 specifications at **intel.com**.
- 2. Update the IPMI channel privileges.

```
racadm set iDRAC.IPMILan.PrivLimit <level>
```

Parameter	Privilege level
<level> = 2</level>	User
<level> = 3</level>	Operator
<level> = 4</level>	Administrator

3. Set the IPMI LAN channel encryption key , if required.

racadm set iDRAC.IPMILan.EncryptionKey <key>

Parameter	Description	
<key></key>	20-character encryption key in a valid hexadecimal format.	

(i) NOTE: The iDRAC IPMI supports the RMCP+ protocol. For more information, see the IPMI 2.0 specifications at intel.com.

Enabling or disabling remote RACADM

You can enable or disable remote RACADM using the iDRAC Web interface or RACADM. You can run up to five remote RACADM sessions in parallel.

(i) NOTE: Remote RACADM is enabled by default.

Enabling or disabling remote RACADM using web interface

- 1. In iDRAC Web interface, go to **iDRAC Settings** > **Services**.
- Under Remote RACADM, select the desired option and click Apply. The remote RACADM is enabled or disabled based on the selection.

Enabling or disabling remote RACADM using RACADM

(i) NOTE: It is recommended to run these commands using local RACADM or firmware RACADM.

• To disable remote RACADM:

racadm set iDRAC.Racadm.Enable 0

To enable remote RACADM:

racadm set iDRAC.Racadm.Enable 1

Disabling local RACADM

The local RACADM is enabled by default. To disable, see Disabling access to modify iDRAC configuration settings on host system.

Enabling IPMI on managed system

On a managed system, use the Dell Open Manage Server Administrator to enable or disable IPMI. For more information, see the *OpenManage Server Administrator User's Guide* available at https://www.dell.com/openmanagemanuals.

(i) NOTE: From iDRAC v2.30.30.30 or later, IPMI supports IPv6 address protocol for Linux-based operating systems.

Setting up iDRAC communication 135

Inserido ao protocolo **22.951.206-4** por: **Pedro Henrique Golin Linhares** em: 31/03/2025 12:49. A autenticidade deste documento pode ser validada no endereço: https://www.eprotocolo.pr.gov.br/spiweb/validarDocumento com o código: fc1c44dfe0686d5547ffd691199be699.



Configuring Linux for serial console during boot in RHEL 6

The following steps are specific to the Linux GRand Unified Bootloader (GRUB). Similar changes are required if a different boot loader is used.

() NOTE: When you configure the client VT100 emulation window, set the window or application that is displaying the redirected Virtual Console to 25 rows x 80 columns to make sure the correct text displays. Else, some text screens may be garbled.

Edit the /etc/grub.conf file as follows:

1. Locate the General Setting sections in the file and add the following:

```
serial --unit=1 --speed=57600 terminal --timeout=10 serial
```

2. Append two options to the kernel line:

kernel console=ttyS1,115200n8r console=tty1

3. Disable GRUB's graphical interface and use the text-based interface. Else, the GRUB screen is not displayed in RAC Virtual Console. To disable the graphical interface, comment-out the line starting with splashimage.

The following example provides a sample /etc/grub.conf file that shows the changes described in this procedure.

```
# grub.conf generated by anaconda
 Note that you do not have to rerun grub after making changes to this file
#
# NOTICE: You do not have a /boot partition. This means that all
# kernel and initrd paths are relative to /, e.g.
#
 root (hd0,0)
 kernel /boot/vmlinuz-version ro root=/dev/sdal
# initrd /boot/initrd-version.img
#boot=/dev/sda
default=0
timeout=10
#splashimage=(hd0,2)/grub/splash.xpm.gz
serial --unit=1 --speed=57600
terminal --timeout=10 serial
title Red Hat Linux Advanced Server (2.4.9-e.3smp) root (hd0,0)
kernel /boot/vmlinuz-2.4.9-e.3smp ro root=/dev/sda1 hda=ide-scsi console=ttyS0
console=ttyS1,115200n8r
initrd /boot/initrd-2.4.9-e.3smp.img
title Red Hat Linux Advanced Server-up (2.4.9-e.3) root (hd0,00)
kernel /boot/vmlinuz-2.4.9-e.3 ro root=/dev/sda1 s
initrd /boot/initrd-2.4.9-e.3.im
```

4. To enable multiple GRUB options to start Virtual Console sessions through the RAC serial connection, add the following line to all options:

console=ttyS1,115200n8r console=tty1

The example shows console=ttyS1, 57600 added to the first option.

(i) NOTE: If the boot loader or operating system provides serial redirection such as GRUB or Linux, then the BIOS Redirection After Boot setting must be disabled. This is to avoid potential race condition of multiple components accessing the serial port.

Enabling login to the virtual console after boot

In the file /etc/inittab, add a new line to configure agetty on the COM2 serial port:



co:2345:respawn:/sbin/agetty -h -L 57600 ttyS1 ansi

The following example shows a sample file with the new line.

#inittab This file describes how the INIT process should set up #the system in a certain run-level. #Author:Miquel van Smoorenburg #Modified for RHS Linux by Marc Ewing and Donnie Barnes #Default runlevel. The runlevels used by RHS are: #0 - halt (Do NOT set initdefault to this) #1 - Single user mode #2 - Multiuser, without NFS (The same as 3, if you do not have #networking) #3 - Full multiuser mode #4 - unused #5 - X11 #6 - reboot (Do NOT set initdefault to this) id:3:initdefault: #System initialization. si::sysinit:/etc/rc.d/rc.sysinit 10:0:wait:/etc/rc.d/rc 0 l1:1:wait:/etc/rc.d/rc 1 l2:2:wait:/etc/rc.d/rc 13:3:wait:/etc/rc.d/rc 3 14:4:wait:/etc/rc.d/rc 4 15:5:wait:/etc/rc.d/rc 5 l6:6:wait:/etc/rc.d/rc 6 #Things to run in every runlevel. ud::once:/sbin/update ud::once:/sbin/update #Trap CTRL-ALT-DELETE ca::ctrlaltdel:/sbin/shutdown -t3 -r now #When our UPS tells us power has failed, assume we have a few #minutes of power left. Schedule a shutdown for 2 minutes from now. #This does, of course, assume you have power installed and your #UPS is connected and working correctly. pf::powerfail:/sbin/shutdown -f -h +2 "Power Failure; System Shutting Down" #If power was restored before the shutdown kicked in, cancel it. pr:12345:powerokwait:/sbin/shutdown -c "Power Restored; Shutdown Cancelled" #Run gettys in standard runlevels co:2345:respawn:/sbin/agetty -h -L 57600 ttyS1 ansi 1:2345:respawn:/sbin/mingetty tty1 2:2345:respawn:/sbin/mingetty tty2

3:2345:respawn:/sbin/mingetty tty3 4:2345:respawn:/sbin/mingetty tty4 5:2345:respawn:/sbin/mingetty tty5 6:2345:respawn:/sbin/mingetty tty6 #Run xdm in runlevel 5 #xdm is now a separate service x:5:respawn:/etc/X11/prefdm -nodaemon

In the file **/etc/securetty** add a new line with the name of the serial tty for COM2:

ttvS1

The following example shows a sample file with the new line.

NOTE: Use the Break Key Sequence (~B) to execute the Linux Magic SysRq key commands on serial console using IPMI Tool.

vc/1 vc/2 vc/3 vc/4 vc/5

vc/6

- vc/7 vc/8
- vc/9

	Fis. <u>1180</u>
	Z Mov. 33
	(Q) (5)
VC/10	400 DO 1
VC/11	
ttyl	
tty2	
tty3	
tty4	
tty5	
tty6	
ttv7	
tty8	
tty9	
ttv10	
ttv11	
++ 1/21	

Configuring serial terminal in RHEL 7

To configure serial terminal in RHEL 7:

1. Add, or update the following lines to /etc/default/grub:

GRUB_CMDLINE_LINUX_DEFAULT="console=tty0 console=ttyS0,115200n8"

GRUB_TERMINAL="console serial"

GRUB_SERIAL_COMMAND="serial --speed=115200 --unit=0 --word=8 --parity=no --stop=1"

GRUB_CMDLINE_LINUX_DEFAULT applies this configuration only to the default menu entry, use GRUB_CMDLINE_LINUX to apply it to all the menu entries.

Each line should only appear once within the /etc/default/grub. If the line already exists, then modify it to avoid another copy. Therefore, only one GRUB CMDLINE LINUX DEFAULT line is allowed.

- 2. Rebuild the /boot/grub.cfg configuration file by running the grub2-mkconfig -o command as follows:
 - on BIOS-based systems:

```
~]# grub2-mkconfig -o /boot/grub2/grub.cfg
```

• on UEFI-based systems:

~] # grub2-mkconfig -o /boot/efi/EFI/redhat/grub.cfg

For more information, see the RHEL 7 System Administrator's Guide at redhat.com.

Controlling GRUB from serial console

You can configure GRUB to use the serial console instead of the VGA console. This allows you to interrupt the boot process and choose a different kernel or add kernel parameters, for example, to boot into single user mode.



To configure GRUB to use serial console, comment out the splash image and add the serial and terminal options to grub.conf:

[root@localhost ~]# cat /boot/grub/grub.conf		
# grub.conf generated by anaconda		
#		
# Note that you do not have to rerun grub after making changes to this	file	
# NOTICE: You have a /boot partition. This means that		
# all kernel and initrd paths are relative to /boot/, eg.		
# root (hd0,0)		
<pre># kernel /vmlinuz-version ro root=/dev/hda2</pre>		
<pre># initrd /initrd-version.img</pre>		
#boot=/dev/hda		
default=0		
timeout=10		
<pre>#splashimage=(hd0,0)/grub/splash.xpm.gz</pre>		
serialunit=0speed=1152001		

NOTE: Restart the system for the settings to take effect.

Supported SSH cryptography schemes

To communicate with iDRAC using SSH protocol, it supports multiple cryptography schemes listed in the following table.

Table 19. SSH cryptography schemes

Scheme Type	Algorithms
Asymmetric Cryptography	
Public key	ssh-rsa
	ecdsa-sha2-nistp256
Symmetric Cryptography	
Key Exchange	curve25519-sha256@libssh.org
	ecdh-sha2-nistp256
	ecdh-sha2-nistp384
	ecdh-sha2-nistp521
	diffie-hellman-group-exchange-sha256

Setting up iDRAC communication 139



Table 19. SSH cryptography schemes (continued)

Scheme Type	Algorithms
	diffie-hellman-group14-sha1
Encryption	chacha20-poly1305@openssh.com aes128-ctr aes192-ctr aes256-ctr aes128-gcm@openssh.com aes256-gcm@openssh.com
MAC	hmac-sha1 hmac-ripemd160 umac-64@openssh.com
Compression	INONE

NOTE: If you enable OpenSSH 7.0 or later, DSA public key support is disabled. To ensure better security for iDRAC, Dell recommends not enabling DSA public key support.

Using public key authentication for SSH

iDRAC supports the Public Key Authentication (PKA) over SSH. This is a licensed feature. When the PKA over SSH is set up and used correctly, you must enter the user name while logging into iDRAC. This is useful for setting up automated scripts that perform various functions. The uploaded keys must be in RFC 4716 or OpenSSH format. Else, you must convert the keys into that format.

In any scenario, a pair of private and public key must be generated on the management station. The public key is uploaded to iDRAC local user and private key is used by the SSH client to establish the trust relationship between the management station and iDRAC.

You can generate the public or private key pair using:

- PuTTY Key Generator application for clients running Windows
- ssh-keygen CLI for clients running Linux.

CAUTION: This privilege is normally reserved for users who are members of the Administrator user group on iDRAC. However, users in the 'Custom' user group can be assigned this privilege. A user with this privilege can modify any user's configuration. This includes creation or deletion of any user, SSH Key management for users, and so on. For these reasons, assign this privilege carefully.

CAUTION: The capability to upload, view, and/ or delete SSH keys is based on the 'Configure Users' user privilege. This privilege allows user(s) to configure another user's SSH key. You should grant this privilege carefully.

Generating public keys for Windows

To use the PuTTY Key Generator application to create the basic key:

- 1. Start the application and select RSA for the key type.
- 2. Enter the number of bits for the key. The number of bits must be between 2048 and 4096 bits.
- Click Generate and move the mouse in the window as directed. The keys are generated.
- 4. You can modify the key comment field.
- 5. Enter a passphrase to secure the key.
- 6. Save the public and private key.



Generating public keys for Linux

To use the *ssh-keygen* application to create the basic key, open a terminal window and at the shell prompt, enter *ssh-keygen* -t *rsa* -b 2048 -C testing

where:

- -t is rsa.
- -b specifies the bit encryption size between 2048 and 4096.
- -C allows modifying the public key comment and is optional.

(i) NOTE: The options are case-sensitive.

Follow the instructions. After the command executes, upload the public file.

CAUTION: Keys generated from the Linux management station using ssh-keygen are in non-4716 format. Convert the keys into the 4716 format using ssh-keygen -e -f /root/.ssh/id_rsa.pub > std_rsa.pub. Do not change the permissions of the key file. The conversion must be done using default permissions.

(i) NOTE: iDRAC does not support ssh-agent forward of keys.

Uploading SSH keys

You can upload up to four public keys *per user* to use over an SSH interface. Before adding the public keys, make sure that you view the keys if they are set up, so that a key is not accidentally overwritten.

When adding new public keys, make sure that the existing keys are not at the index where the new key is added. iDRAC does not perform checks to make sure previous key(s) are deleted before a new key(s) are added. When a new key is added, it is usable if the SSH interface is enabled.

Uploading SSH keys using web interface

To upload the SSH keys:

- In the iDRAC Web interface, go to iDRAC Settings > Users > Local Users. The Local Users page is displayed.
- 2. In the User ID column, click a user ID number. The Users Main Menu page is displayed.
- **3.** Under SSH Key Configurations, select Upload SSH Key(s) and click Next. The Upload SSH Key(s) page is displayed.
- 4. Upload the SSH keys in one of the following ways:
 - Upload the key file.
 - Copy the contents of the key file into the text box

For more information, see iDRAC Online Help.

5. Click Apply.

Uploading SSH keys using RACADM

To upload the SSH keys, run the following command:

(i) NOTE: You cannot upload and copy a key at the same time.

- For local RACADM: racadm sshpkauth -i <2 to 16> -k <1 to 4> -f <filename>
- From remote RACADM using or SSH: racadm sshpkauth -i <2 to 16> -k <1 to 4> -t <key-text>

For example, to upload a valid key to iDRAC User ID 2 in the first key space using a file, run the following command:

\$ racadm sshpkauth -i 2 -k 1 -f pkkey.key

(i) **NOTE:** The -f option is not supported on ssh/serial RACADM.



Viewing SSH keys

You can view the keys that are uploaded to iDRAC.

Viewing SSH keys using web interface

To view the SSH keys:

- In Web interface, go to iDRAC Settings > Users. The Local Users page is displayed.
- 2. In the User ID column, click a user ID number. The Users Main Menu page is displayed.
- **3.** Under SSH Key Configurations, select View/Remove SSH Key(s) and click Next. The View/Remove SSH Key(s) page is displayed with the key details.

Deleting SSH keys

Before deleting the public keys, make sure that you view the keys if they are set up, so that a key is not accidentally deleted.

Deleting SSH keys using web interface

To delete the SSH key(s):

- 1. In Web interface, go to **iDRAC Settings** > **Users**. The **Local Users** page is displayed.
- In the ID column, select a user ID number, click Edit. The Edit User page is displayed.
- **3.** Under **SSH Key Configurations**, select a SSH Key and click **Edit**. The **SSH Key** page displays the **Edit From** details.
- Select Remove for the key(s) you want to delete, and click Apply. The selected key(s) is deleted.

Deleting SSH keys using RACADM

To delete the SSH key(s), run the following commands:

- Specific key racadm sshpkauth -i <2 to 16> -d -k <1 to 4>
- All keys racadm sshpkauth -i <2 to 16> -d -k all



Configuring user accounts and privileges

You can setup user accounts with specific privileges (*role-based authority*) to manage your system using iDRAC and maintain system security. By default iDRAC is configured with a local administrator account. The default iDRAC user name and password are provided with the system badge. As an administrator, you can setup user accounts to allow other users to access iDRAC. For more information see the documentation for the server.

You can setup local users or use directory services such as Microsoft Active Directory or LDAP to setup user accounts. Using a directory service provides a central location for managing authorized user accounts.

iDRAC supports role-based access to users with a set of associated privileges. The roles are administrator, operator, read only, or none. The role defines the maximum privileges available.

Topics:

- iDRAC user roles and privileges
- Recommended characters in user names and passwords
- Configuring local users
- Configuring Active Directory users
- Configuring generic LDAP users

iDRAC user roles and privileges

The iDRAC role and privilege names have changed from earlier generation of servers. The role names are:

Table 20. iDRAC roles

Current Generation	Prior Generation	Privileges
Administrator	Administrator	Login, Configure, Configure Users, Logs, System Control, Access Virtual Console, Access Virtual Media, System Operations, Debug
Operator	Power User	Login, Configure, System Control, Access Virtual Console, Access Virtual Media, System Operations, Debug
Read Only	Guest User	Login
None	None	None

The following table describes the user privileges:

Table 21. iDRAC user privileges

Current Generation	Prior Generation	Description
Login	Login to iDRAC	Enables the user to log in to iDRAC.
Configure	Configure iDRAC	Enables the user to configure iDRAC. With this privilege, a user can also configure power management, virtual console, virtual media, licenses, system settings, storage devices, BIOS settings, SCP and so on.
NOTE: The administrator role overrides all the privileges from the other components such as BIOS setup password.		
Configure Users	Configure Users	Enables the user to allow specific users to access the system.
Logs	Clear Logs	Enables the user to clear only the System Event Log (SEL).

Configuring user accounts and privileges 143



Table 21. iDRAC user privileges (continued)

Current Generation	Prior Generation	Description
System Control	Control and configure system	Allows power cycling the host system.
Access Virtual Console	Access Virtual Console Redirection (for blade servers) Access Virtual Console (for rack and tower servers)	Enables the user to run Virtual Console.
Access Virtual Media	Access Virtual Media	Enables the user to run and use Virtual Media.
System Operations	Test Alerts	Allows user initiated and generated events, and information is sent as an asynchronous notification and logged.
Debug	Execute Diagnostic Commands	Enables the user to run diagnostic commands.

Recommended characters in user names and passwords

This section provides details about the recommended characters while creating and using user names and passwords.

(i) NOTE: The password must include one uppercase and one lower case letter, one number and a special character.

Use the following characters while creating user names and passwords:

Table 22. Recommended characters for user names

Characters	Length
0-9	1–16
A-Z	
a-z	
-!#\$%&()*;?[\]^_`{ }~+<=>	

Table 23. Recommended characters for passwords

Characters	Length
0-9	1-40
A-Z	
a-z	
'-!"#\$%&()*,./:;?@[\]^_`{ }~+<=>	

NOTE: You may be able to create user names and passwords that include other characters. However, to ensure compatibility with all interfaces, Dell recommends using only the characters listed here.

NOTE: The characters allowed in user names and passwords for network shares are determined by the network-share type. iDRAC supports valid characters for network share credentials as defined by the share type, except <, >, and , (comma).

() NOTE: To improve security, it is recommended to use complex passwords that have eight or more characters and include lowercase alphabets, uppercase alphabets, numbers, and special characters. It is also recommended to regularly change the passwords, if possible.

144 Configuring user accounts and privileges



Configuring local users

You can configure up to 16 local users in iDRAC with specific access permissions. Before you create an iDRAC user, verify if any current users exist. You can set user names, passwords, and roles with the privileges for these users. The user names and passwords can be changed using any of the iDRAC secured interfaces (that is, web interface, RACADM or WSMan). You can also enable or disable SNMPv3 authentication for each user.

Configuring local users using iDRAC web interface

To add and configure local iDRAC users:

(i) NOTE: You must have Configure Users permission to create an iDRAC user.

- In the iDRAC Web interface, go to iDRAC Settings > User. The Local Users page is displayed.
- 2. In the User ID column, select a user ID number and click Edit.

(i) NOTE: User 1 is reserved for the IPMI anonymous user and you cannot change this configuration.

The User Configuration page is displayed.

3. Add User Account Settings and Advanced Settings details to configure the user account.

() **NOTE:** Enable the user ID and specify the user name, password, and user role (access privileges) for the user. You can also enable LAN privilege level, Serial port privilege level, serial over LAN status, SNMPv3 authentication, authentication type and the privacy type for the user. For more information about the options, see the *iDRAC Online Help*.

4. Click **Save**. The user is created with the required privileges.

Configuring local users using RACADM

(i) NOTE: You must be logged in as user **root** to execute RACADM commands on a remote Linux system.

You can configure single or multiple iDRAC users using RACADM.

To configure multiple iDRAC users with identical configuration settings, follow these procedures:

- Use the RACADM examples in this section as a guide to create a batch file of RACADM commands and then execute the batch file on each managed system.
- Create the iDRAC configuration file and execute the racadm set command on each managed system using the same configuration file.

If you are configuring a new iDRAC or if you have used the racadm racresetcfg command, then check for the default iDRAC user name and password on the system badge. The racadm racresetcfg command resets the iDRAC to the default values.

NOTE: If SEKM is enabled on the server, then disable SEKM using the racadm sekm disable command before using this command. This can avoid any storage devices being locked out which are secured by iDRAC, if SEKM settings are erased from iDRAC by executing this command.

(i) NOTE: Users can be enabled and disabled over time. As a result, a user may have a different index number on each iDRAC.

To verify if a user exists, type the following command once for each index (1–16):

racadm get iDRAC.Users.<index>.UserName

Several parameters and object IDs are displayed with their current values. The key field is iDRAC.Users.UserName=. If a user name is displayed after =, that index number is taken.

(i) NOTE: You can utilize

racadm get -f <myfile.cfg>

Fis.1188 Mov. 33

and view or edit the

myfile.cfg

file, which includes all iDRAC configuration parameters.

To enable SNMP v3 authentication for a user, use **SNMPv3AuthenticationType**, **SNMPv3Enable**, **SNMPv3PrivacyType** objects. For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

If you use the Server Configuration Profile file to configure users, use the **AuthenticationProtocol**, **ProtocolEnable**, and **PrivacyProtocol** attributes to enable SNMPv3 authentication.

Adding iDRAC user using RACADM

1. Set the index and user name.

racadm set idrac.users.<index>.username <user name>

Parameter	Description
<index></index>	Unique index of the user
<user_name></user_name>	User name

2. Set the password.

racadm set idrac.users.<index>.password <password>

3. Set the user privileges.

For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

4. Enable the user.

racadm set.idrac.users.<index>.enable 1

To verify, use the following command:

racadm get idrac.users.<index>

For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Enabling iDRAC user with permissions

To enable a user with specific administrative permissions (role-based authority):

1. Locate an available user index.

racadm get iDRAC.Users <index>

2. Type the following commands with the new user name and password.

racadm set iDRAC.Users.<index>.Privilege <user privilege bit mask value>

NOTE: The default privilege value is 0, which indicates the user has no privileges enabled. For a list of valid bit-mask values for specific user privileges, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.



Configuring Active Directory users

If your company uses the Microsoft Active Directory software, you can configure the software to provide access to iDRAC, allowing you to add and control iDRAC user privileges to your existing users in your directory service. This is a licensed feature.

You can configure user authentication through Active Directory to log in to the iDRAC. You can also provide role-based authority, which enables an administrator to configure specific privileges for each user.

NOTE: For any deployment done via MX Template and CA validation is enabled within template, the user must upload CA certificates at first login or before changing Authentication Service from LDAP to Active Directory or vice versa.

Prerequisites for using Active Directory authentication for iDRAC

To use the Active Directory authentication feature of iDRAC, make sure that you have:

- Deployed an Active Directory infrastructure. See the Microsoft website for more information.
- Integrated PKI into the Active Directory infrastructure. iDRAC uses the standard Public Key Infrastructure (PKI) mechanism to authenticate securely into the Active Directory. See the Microsoft website for more information.
- Enabled the Secure Socket Layer (SSL) on all domain controllers that iDRAC connects to for authenticating to all the domain controllers.

Enabling SSL on domain controller

When iDRAC authenticates users with an Active Directory domain controller, it starts an SSL session with the domain controller. At this time, the domain controller must publish a certificate signed by the Certificate Authority (CA)—the root certificate of which is also uploaded into iDRAC. For iDRAC to authenticate to *any* domain controller—whether it is the root or the child domain controller—that domain controller must have an SSL-enabled certificate signed by the domain's CA.

If you are using Microsoft Enterprise Root CA to automatically assign all your domain controllers to an SSL certificate, you must:

- 1. Install the SSL certificate on each domain controller.
- 2. Export the Domain Controller Root CA Certificate to iDRAC.
- 3. Import iDRAC Firmware SSL Certificate.

Installing SSL certificate for each domain controller

To install the SSL certificate for each controller:

- 1. Click Start > Administrative Tools > Domain Security Policy.
- 2. Expand the Public Key Policies folder, right-click Automatic Certificate Request Settings and click Automatic Certificate Request.
 - The Automatic Certificate Request Setup Wizard is displayed.
- 3. Click Next and select Domain Controller.
- 4. Click Next and click Finish. The SSL certificate is installed.

Exporting domain controller root CA certificate to iDRAC

To export the domain controller root CA certificate to iDRAC:

- 1. Locate the domain controller that is running the Microsoft Enterprise CA service.
- 2. Click Start > Run.
- 3. Enter mmc and click OK.
- 4. In the Console 1 (MMC) window, click File (or Console) and select Add/Remove Snap-in.
- 5. In the Add/Remove Snap-In window, click Add.
- 6. In the Standalone Snap-In window, select Certificates and click Add.
- 7. Select Computer and click Next.
- 8. Select Local Computer, click Finish, and click OK.
- 9. In the Console 1 window, go to Certificates Personal Certificates folder.



- 10. Locate and right-click the root CA certificate, select All Tasks, and click Export....
- 11. In the Certificate Export Wizard, click Next, and select No do not export the private key.
- 12. Click Next and select Base-64 encoded X.509 (.cer) as the format.
- 13. Click **Next** and save the certificate to a directory on your system.
- 14. Upload the certificate you saved in step 13 to iDRAC.

Importing iDRAC firmware SSL certificate

iDRAC SSL certificate is the identical certificate used for iDRAC Web server. All iDRAC controllers are shipped with a default self-signed certificate.

If the Active Directory Server is set to authenticate the client during an SSL session initialization phase, you need to upload iDRAC Server certificate to the Active Directory Domain controller. This additional step is not required if the Active Directory does not perform a client authentication during an SSL session's initialization phase.

NOTE: If iDRAC firmware SSL certificate is CA-signed and the certificate of that CA is already in the domain controller's Trusted Root Certificate Authority list, do not perform the steps in this section.

To import iDRAC firmware SSL certificate to all domain controller trusted certificate lists:

1. Download iDRAC SSL certificate using the following RACADM command:

racadm sslcertdownload -t 1 -f <RAC SSL certificate>

- 2. On the domain controller, open an MMC Console window and select Certificates > Trusted Root Certification Authorities.
- 3. Right-click Certificates, select All Tasks and click Import.
- 4. Click Next and browse to the SSL certificate file.
- 5. Install iDRAC SSL Certificate in each domain controller's Trusted Root Certification Authority.

If you have installed your own certificate, make sure that the CA signing your certificate is in the **Trusted Root Certification Authority** list. If the Authority is not in the list, you must install it on all your domain controllers.

- 6. Click **Next** and select whether you want Windows to automatically select the certificate store based on the type of certificate, or browse to a store of your choice.
- 7. Click Finish and click OK. The iDRAC firmware SSL certificate is imported to all domain controller trusted certificate lists.

Supported Active Directory authentication mechanisms

You can use Active Directory to define iDRAC user access using two methods:

- Standard schema solution, which uses Microsoft's default Active Directory group objects only.
- Extended schema solution, which has customized Active Directory objects. All the access control objects are maintained in Active Directory. It provides maximum flexibility to configure user access on different iDRACs with varying privilege levels.

Standard schema Active Directory overview

As shown in the following figure, using standard schema for Active Directory integration requires configuration on both Active Directory and iDRAC.





Figure 1. Configuration of iDRAC with active directory standard schema

In Active Directory, a standard group object is used as a role group. A user who has iDRAC access is a member of the role group. To give this user access to a specific iDRAC, the role group name and its domain name need to be configured on the specific iDRAC. The role and the privilege level are defined on each iDRAC and not in the Active Directory. You can configure up to 15 role groups in each iDRAC. Table reference no shows the default role group privileges.

Role Groups	Default Privilege Level	Permissions Granted	Bit Mask
Role Group 1	None	Log in to iDRAC, Configure iDRAC, Configure Users, Clear Logs, Execute Server Control Commands, Access Virtual Console, Access Virtual Media, Test Alerts, Execute Diagnostic Commands	0x000001ff
Role Group 2	None	Log in to iDRAC, Configure iDRAC, Execute Server Control Commands, Access Virtual Console, Access Virtual Media, Test Alerts, Execute Diagnostic Commands	0x00000f9
Role Group 3	None	Log in to iDRAC	0x0000001
Role Group 4	None	No assigned permissions	0×0000000
Role Group 5	None	No assigned permissions	0x0000000

Table 24. Default role group privileges

(i) NOTE: The Bit Mask values are used only when setting Standard Schema with the RACADM.

Single domain versus multiple domain scenarios

If all the login users and role groups, including the nested groups, are in the same domain, then only the domain controllers' addresses must be configured on iDRAC. In this single domain scenario, any group type is supported.

If all the login users and role groups, or any of the nested groups, are from multiple domains, then Global Catalog server addresses must be configured on iDRAC. In this multiple domain scenario, all the role groups and nested groups, if any, must be a Universal Group type.



Configuring Standard schema Active Directory

Before configuring the standard schema Active Directory, ensure that:

- You have the iDRAC Enterprise or Datacenter license.
- The configuration is performed on a server that is used as the Domain Controller.
- The dat, time and time zone on the server are correct.
- The iDRAC network settings are configured, or in iDRAC web interface go to iDRAC Settings > Connectivity > Network > Common Settings to configure the network settings.

To configure iDRAC for an Active Directory login access:

- 1. On an Active Directory server (domain controller), open the Active Directory Users and Computers Snap-in.
- 2. Create the iDRAC groups and users.
- 3. Configure the group name, domain name, and the role privileges on iDRAC using the iDRAC web interface or RACADM.

Configuring Active Directory with Standard schema using iDRAC web interface

(i) NOTE: For information about the various fields, see the *iDRAC Online Help*.

- In the iDRAC web interface, go to iDRAC Settings > Users > Directory Services. The Directory Service page is displayed.
- Select the Microsoft Active Directory option and then click Edit. The Active Directory Configuration and Management page is displayed.
- 3. Click Configure Active Directory. The Active Directory Configuration and Management Step 1 of 4 page is displayed.
- 4. Optionally, enable certificate validation and upload the CA-signed digital certificate used during initiation of SSL connections when communicating with the Active Directory (AD) server. For this, the Domain Controllers and Global Catalog FQDN must be specified. This is done in the next steps. And hence the DNS should be configured properly in the network settings.
- 5. Click Next.

The Active Directory Configuration and Management Step 2 of 4 page is displayed.

6. Enable Active Directory and specify the location information about Active Directory servers and user accounts. Also, specify the time iDRAC must wait for responses from Active Directory during iDRAC login.

NOTE: If certificate validation is enabled, specify the Domain Controller Server addresses and the Global Catalog FQDN. Make sure that DNS is configured correctly under **iDRAC Settings** > **Network**.

- 7. Click Next. The Active Directory Configuration and Management Step 3 of 4 page is displayed.
- Select Standard Schema and click Next.
 The Active Directory Configuration and Management Step 4a of 4 page is displayed.
- 9. Enter the location of Active Directory global catalog server(s) and specify privilege groups used to authorize users.
- Click a Role Group to configure the control authorization policy for users under the standard schema mode. The Active Directory Configuration and Management Step 4b of 4 page is displayed.
- 11. Specify the privileges and click Apply.

The settings are applied and the Active Directory Configuration and Management Step 4a of 4 page is displayed.

12. Click **Finish**. The Active Directory settings for standard schema are configured.

Configuring Active Directory with Standard schema using RACADM

1. Use the following commands:

```
racadm set iDRAC.ActiveDirectory.Enable 1
racadm set iDRAC.ActiveDirectory.Schema 2
racadm set iDRAC.ADGroup.Name <common name of the role group>
racadm set iDRAC.ADGroup.Domain <fully qualified domain name>
racadm set iDRAC.ADGroup.Privilege <Bit-mask value for specific RoleGroup permissions>
racadm set iDRAC.ActiveDirectory.DomainController1 <fully qualified domain name or IP</pre>
```

150 Configuring user accounts and privileges



```
address of the domain controller>
racadm set iDRAC.ActiveDirectory.DomainController2 <fully qualified domain name or IP
address of the domain controller>
racadm set iDRAC.ActiveDirectory.DomainController3 <fully qualified domain name or IP
address of the domain controller>
racadm set iDRAC.ActiveDirectory.GlobalCatalog1 <fully qualified domain name or IP
address of the domain controller>
racadm set iDRAC.ActiveDirectory.GlobalCatalog2 <fully qualified domain name or IP
address of the domain controller>
racadm set iDRAC.ActiveDirectory.GlobalCatalog2 <fully qualified domain name or IP
address of the domain controller>
racadm set iDRAC.ActiveDirectory.GlobalCatalog3 <fully qualified domain name or IP
address of the domain controller>
```

- Enter the Fully Qualified Domain Name (FQDN) of the domain controller, not the FQDN of the domain. For example, enter servername.dell.com instead of dell.com.
- For bit-mask values for specific Role Group permissions, see Default role group privileges.
- You must provide at least one of the three domain controller addresses. iDRAC attempts to connect to each of the configured addresses one-by-one until it makes a successful connection. With Standard Schema, these are the addresses of the domain controllers where the user accounts and the role groups are located.
- The Global Catalog server is only required for standard schema when the user accounts and role groups are in different domains. In multiple domain case, only the Universal Group can be used.
- If certificate validation is enabled, the FQDN or IP address that you specify in this field must match the Subject or Subject Alternative Name field of your domain controller certificate.
- To disable the certificate validation during SSL handshake, use the following command:

racadm set iDRAC.ActiveDirectory.CertValidationEnable 0

In this case, no Certificate Authority (CA) certificate needs to be uploaded.

To enforce the certificate validation during SSL handshake (optional), use the following command:

racadm set iDRAC.ActiveDirectory.CertValidationEnable 1

In this case, you must upload the CA certificate using the following command:

racadm sslcertupload -t 0x2 -f <ADS root CA certificate>

NOTE: If certificate validation is enabled, specify the Domain Controller Server addresses and the Global Catalog FQDN. Ensure that DNS is configured correctly under **Overview** > **iDRAC Settings** > **Network**.

Using the following RACADM command may be optional.

racadm sslcertdownload -t 1 -f <RAC SSL certificate>

2. If DHCP is enabled on iDRAC and you want to use the DNS provided by the DHCP server, enter the following command:

racadm set iDRAC.IPv4.DNSFromDHCP 1

3. If DHCP is disabled on iDRAC or you want manually enter the DNS IP address, enter the following RACADM command:

```
racadm set iDRAC.IPv4.DNSFromDHCP 0
racadm set iDRAC.IPv4.DNSFromDHCP.DNS1 <primary DNS IP address>
racadm set iDRAC.IPv4.DNSFromDHCP.DNS2 <secondary DNS IP address>
```

4. If you want to configure a list of user domains so that you only need to enter the user name when logging in to the web interface, use the following command:

```
racadm set iDRAC.UserDomain.<index>.Name <fully qualified domain name or IP Address
of the domain controller>
```

You can configure up to 40 user domains with index numbers between 1 and 40.

Extended schema Active Directory overview

Using the extended schema solution requires the Active Directory schema extension.

Configuring user accounts and privileges 151



Best practices for extended schema

The extended schema uses Dell association objects to join iDRAC and permission. This allows you to use iDRAC based on the overall permissions granted. The default Access Control List (ACL) of Dell Association objects allows Self and Domain Administrators to manage the permissions and scope of iDRAC objects.

By default, the Dell Association objects do not inherit all permissions from the parent Active Directory objects. If you enable inheritance for the Dell Association object, the inherited permissions for that association object are granted to the selected users and groups. This may result in unintended privileges being provided to the iDRAC.

To use the Extended Schema securely, Dell recommends not enabling inheritance on Dell Association objects within the extended schema implementation.

Active directory schema extensions

The Active Directory data is a distributed database of *attributes* and *classes*. The Active Directory schema includes the rules that determine the type of data that can be added or included in the database. The user class is one example of a *class* that is stored in the database. Some example user class attributes can include the user's first name, last name, phone number, and so on. You can extend the Active Directory database by adding your own unique *attributes* and *classes* for specific requirements. Dell has extended the schema to include the necessary changes to support remote management authentication and authorization using Active Directory.

Each *attribute* or *class* that is added to an existing Active Directory Schema must be defined with a unique ID. To maintain unique IDs across the industry, Microsoft maintains a database of Active Directory Object Identifiers (OIDs) so that when companies add extensions to the schema, they can be guaranteed to be unique and not to conflict with each other. To extend the schema in Microsoft's Active Directory, Dell received unique OIDs, unique name extensions, and uniquely linked attribute IDs for the attributes and classes that are added into the directory service:

- Extension is: dell
- Base OID is: 1.2.840.113556.1.8000.1280
- RAC LinkID range is: 12070 to 12079

Overview of iDRAC schema extensions

Dell has extended the schema to include an *Association, Device,* and *Privilege* property. The *Association* property is used to link together the users or groups with a specific set of privileges to one or more iDRAC devices. This model provides an administrator maximum flexibility over the different combinations of users, iDRAC privileges, and iDRAC devices on the network without much complexity.

For each physical iDRAC device on the network that you want to integrate with Active Directory for authentication and authorization, create at least one association object and one iDRAC device object. You can create multiple association objects, and each association object can be linked to as many users, groups of users, or iDRAC device objects as required. The users and iDRAC user groups can be members of any domain in the enterprise.

However, each association object can be linked (or, may link users, groups of users, or iDRAC device objects) to only one privilege object. This example allows an administrator to control each user's privileges on specific iDRAC devices.

iDRAC device object is the link to iDRAC firmware for querying Active Directory for authentication and authorization. When iDRAC is added to the network, the administrator must configure iDRAC and its device object with its Active Directory name so that users can perform authentication and authorization with Active Directory. Additionally, the administrator must add iDRAC to at least one association object for users to authenticate.

The following figure shows that the association object provides the connection that is needed for the authentication and authorization.





Figure 2. Typical setup for active directory objects

You can create as many or as few association objects as required. However, you must create at least one Association Object, and you must have one iDRAC Device Object for each iDRAC device on the network that you want to integrate with Active Directory for Authentication and Authorization with iDRAC.

The Association Object allows for as many or as few users and/or groups as well as iDRAC Device Objects. However, the Association Object only includes one Privilege Object per Association Object. The Association Object connects the Users who have Privileges on iDRAC devices.

The Dell extension to the ADUC MMC Snap-in only allows associating the Privilege Object and iDRAC Objects from the same domain with the Association Object. The Dell extension does not allow a group or an iDRAC object from other domains to be added as a product member of the Association Object.

When adding Universal Groups from separate domains, create an Association Object with Universal Scope. The Default Association objects created by the Dell Schema Extender Utility are Domain Local Groups and they do not work with Universal Groups from other domains.

Users, user groups, or nested user groups from any domain can be added into the Association Object. Extended Schema solutions support any user group type and any user group nesting across multiple domains allowed by Microsoft Active Directory.

Accumulating privileges using Extended Schema

The Extended Schema Authentication mechanism supports Privilege Accumulation from different privilege objects associated with the same user through different Association Objects. In other words, Extended Schema Authentication accumulates privileges to allow the user the super set of all assigned privileges corresponding to the different privilege objects associated with the same user.

The following figure provides an example of accumulating privileges using Extended Schema.



Figure 3. Privilege accumulation for a user



The figure shows two Association Objects—A01 and A02. User1 is associated to iDRAC2 through both association objects.

Extended Schema Authentication accumulates privileges to allow the user the maximum set of privileges possible considering the assigned privileges of the different privilege objects associated to the same user.

In this example, User1 has both Priv1 and Priv2 privileges on iDRAC2. User1 has Priv1 privileges on iDRAC1 only. User2 has Priv1 privileges on both iDRAC1 and iDRAC2. In addition, this figure shows that User1 can be in a different domain and can be a member of a group.

Configuring Extended schema Active Directory

To configure Active Directory to access iDRAC:

- 1. Extend the Active Directory schema.
- 2. Extend the Active Directory Users and Computers Snap-in.
- **3.** Add iDRAC users and their privileges to Active Directory.
- 4. Configure iDRAC Active Directory properties using iDRAC Web interface or RACADM.

Extending Active Directory schema

Extending your Active Directory schema adds a Dell organizational unit, schema classes and attributes, and example privileges and association objects to the Active Directory schema. Before you extend the schema, make sure that you have the Schema Admin privileges on the Schema Master FSMO-Role-Owner of the domain forest.

NOTE: The schema extension for this product is different from the previous generations. The earlier schema does not work with this product.

(i) NOTE: Extending the new schema has no impact on previous versions of the product.

You can extend your schema using one of the following methods:

- Dell Schema Extender utility
- LDIF script file

If you use the LDIF script file, the Dell organizational unit is not added to the schema.

The LDIF files and Dell Schema Extender are on your *Dell Systems Management Tools and Documentation* DVD in the following respective directories:

- DVDdrive:\SYSMGMT\ManagementStation\support\OMActiveDirectory_Tools\Remote_Management_Adv anced\LDIF_Files
- OVDdrive>:

\SYSMGMT\ManagementStation\support\OMActiveDirectory_Tools\Remote_Management_Advanced\Sch ema Extender

To use the LDIF files, see the instructions in the readme included in the LDIF_Files directory.

You can copy and run the Schema Extender or LDIF files from any location.

Using Dell Schema Extender

CAUTION: The Dell Schema Extender uses the SchemaExtenderOem.ini file. To make sure that the Dell Schema Extender utility functions properly, do not modify the name of this file.

- 1. In the Welcome screen, click Next.
- 2. Read and understand the warning and click Next.
- 3. Select Use Current Log In Credentials or enter a user name and password with schema administrator rights.
- 4. Click Next to run the Dell Schema Extender.
- 5. Click Finish.

The schema is extended. To verify the schema extension, use the MMC and the Active Directory Schema Snap-in to verify that the Classes and attributes exist. See the Microsoft documentation for details about using the MMC and the Active Directory Schema Snap-in.



Classes and attributes

Table 25. Class definitions for classes added to the active directory schema

Class Name	Assigned Object Identification Number (OID)	
delliDRACDevice	1.2.840.113556.1.8000.1280.1.7.1.1	
delliDRACAssociation	1.2.840.113556.1.8000.1280.1.7.1.2	
dellRAC4Privileges	1.2.840.113556.1.8000.1280.1.1.1.3	
dellPrivileges	1.2.840.113556.1.8000.1280.1.1.1.4	
dellProduct	1.2.840.113556.1.8000.1280.1.1.1.5	

Table 26. DelliDRACdevice class

OID	1.2.840.113556.1.8000.1280.1.7.1.1
Description	Represents the Dell iDRAC device. iDRAC must be configured as delliDRACDevice in Active Directory. This configuration enables iDRAC to send Lightweight Directory Access Protocol (LDAP) queries to Active Directory.
Class Type	Structural Class
SuperClasses	dellProduct
Attributes	dellSchemaVersion dellRacType

Table 27. delliDRACAssociationObject class

OID	1.2.840.113556.1.8000.1280.1.7.1.2	
Description	Represents the Dell Association Object. The Association Object provides the connection between the users and the devices.	
Class Type	Structural Class	
SuperClasses	Group	
Attributes	dellProductMembers dellPrivilegeMember	

Table 28. dellRAC4Privileges class

OID	1.2.840.113556.1.8000.1280.1.1.1.3	
Description	Defines the privileges (Authorization Rights) for iDRAC	
Class Type	Auxiliary Class	
SuperClasses	None	
Attributes	delllsLoginUser	
	dellIsCardConfigAdmin	
	dellIsUserConfigAdmin	



Table 28. dellRAC4Privileges class (continued)

OID	1.2.840.113556.1.8000.1280.1.1.1.3	
dellIsLogClearAdmin		
	dellIsServerResetUser	
	dellIsConsoleRedirectUser	
	dellIsVirtualMediaUser	
	dellIsTestAlertUser	
	dellIsDebugCommandAdmin	

Table 29. dellPrivileges class

OID	1.2.840.113556.1.8000.1280.1.1.1.4
Description	Used as a container Class for the Dell Privileges (Authorization Rights).
Class Type	Structural Class
SuperClasses	User
Attributes	dellRAC4Privileges

Table 30. dellProduct class

OID	1.2.840.113556.1.8000.1280.1.1.1.5
Description	The main class from which all Dell products are derived.
Class Type	Structural Class
SuperClasses	Computer
Attributes	dellAssociationMembers

Table 31. List of attributes added to the active directory schema

Attribute Name/Description	Assigned OID/Syntax Object Identifier	Single Valued
dellPrivilegeMember	1.2.840.113556.1.8000.1280.1.1.2.1	FALSE
List of dellPrivilege Objects that belong to this Attribute.	Distinguished Name (LDAPTYPE_DN 1.3.6.1.4.1.1466.115.121.1.12)	
dellProductMembers	1.2.840.113556.1.8000.1280.1.1.2.2	FALSE
List of dellRacDevice and DelliDRACDevice Objects that belong to this role. This attribute is the forward link to the dellAssociationMembers backward link. Link ID: 12070	Distinguished Name (LDAPTYPE_DN 1.3.6.1.4.1.1466.115.121.1.12)	
dellIsLoginUser	1.2.840.113556.1.8000.1280.1.1.2.3	TRUE
TRUE if the user has Login rights on the device.	Boolean (LDAPTYPE_BOOLEAN 1.3.6.1.4.1.1466.115.121.1.7)	
dellIsCardConfigAdmin	1.2.840.113556.1.8000.1280.1.1.2.4	TRUE

Inserido ao protocolo **22.951.206-4** por: **Pedro Henrique Golin Linhares** em: 31/03/2025 12:49. A autenticidade deste documento pode ser validada no endereço: https://www.eprotocolo.pr.gov.br/spiweb/validarDocumento com o código: fc1c44dfe0686d5547ffd691199be699.



Table 31. List of attributes added to the active directory schema (continued)

Attribute Name/Description	Assigned OID/Syntax Object Identifier	Single Valued
TRUE if the user has Card Configuration rights on the device.	Boolean (LDAPTYPE_BOOLEAN 1.3.6.1.4.1.1466.115.121.1.7)	
dellIsUserConfigAdmin	1.2.840.113556.1.8000.1280.1.1.2.5	TRUE
TRUE if the user has User Configuration rights on the device.	Boolean (LDAPTYPE_BOOLEAN 1.3.6.1.4.1.1466.115.121.1.7)	
dellsLogClearAdmin	1.2.840.113556.1.8000.1280.1.1.2.6	TRUE
TRUE if the user has Log Clearing rights on the device.	Boolean (LDAPTYPE_BOOLEAN 1.3.6.1.4.1.1466.115.121.1.7)	
dellIsServerResetUser	1.2.840.113556.1.8000.1280.1.1.2.7	TRUE
TRUE if the user has Server Reset rights on the device.	Boolean (LDAPTYPE_BOOLEAN 1.3.6.1.4.1.1466.115.121.1.7)	
dellIsConsoleRedirectUser	1.2.840.113556.1.8000.1280.1.1.2.8	TRUE
TRUE if the user has Virtual Console rights on the device.	Boolean (LDAPTYPE_BOOLEAN 1.3.6.1.4.1.1466.115.121.1.7)	
dellIsVirtualMediaUser	1.2.840.113556.1.8000.1280.1.1.2.9	TRUE
TRUE if the user has Virtual Media rights on the device.	Boolean (LDAPTYPE_BOOLEAN 1.3.6.1.4.1.1466.115.121.1.7)	
dellIsTestAlertUser	1.2.840.113556.1.8000.1280.1.1.2.10	TRUE
TRUE if the user has Test Alert User rights on the device.	Boolean (LDAPTYPE_BOOLEAN 1.3.6.1.4.1.1466.115.121.1.7)	
dellIsDebugCommandAdmin	1.2.840.113556.1.8000.1280.1.1.2.11	TRUE
TRUE if the user has Debug Command Admin rights on the device.	Boolean (LDAPTYPE_BOOLEAN 1.3.6.1.4.1.1466.115.121.1.7)	
dellSchemaVersion	1.2.840.113556.1.8000.1280.1.1.2.12	TRUE
The Current Schema Version is used to update the schema.	Case Ignore String (LDAPTYPE_CASEIGNORESTRING 1.2.840.113556.1.4.905)	
dellRacType	1.2.840.113556.1.8000.1280.1.1.2.13	TRUE
This attribute is the Current RAC Type for the delliDRACDevice object and the backward link to the dellAssociationObjectMembers forward link.	Case Ignore String (LDAPTYPE_CASEIGNORESTRING 1.2.840.113556.1.4.905)	
dellAssociationMembers	1.2.840.113556.1.8000.1280.1.1.2.14	FALSE
List of dellAssociationObjectMembers that belong to this Product. This attribute is the backward link to the dellProductMembers linked attribute.	Distinguished Name (LDAPTYPE_DN 1.3.6.1.4.1.1466.115.121.1.12)	
Link ID: 12071		



Installing Dell extension to the Active Directory users and computers snap-ir

When you extend the schema in Active Directory, you must also extend the Active Directory Users and Computers Snap-in so the administrator can manage iDRAC devices, users and user groups, iDRAC associations, and iDRAC privileges.

When you install your systems management software using the *Dell Systems Management Tools and Documentation* DVD, you can extend the Snap-in by selecting the **Active Directory Users and Computers Snap-in** option during the installation procedure. See the Dell OpenManage Software Quick Installation Guide for additional instructions about installing systems management software. For 64-bit Windows Operating Systems, the Snap-in installer is located under:

<DVDdrive>:\SYSMGMT\ManagementStation\support\OMActiveDirectory_SnapIn64

For more information about the Active Directory Users and Computers Snap-in, see Microsoft documentation.

Adding iDRAC users and privileges to Active Directory

Using the Dell-extended Active Directory Users and Computers Snap-in, you can add iDRAC users and privileges by creating device, association, and privilege objects. To add each object, perform the following:

- Create an iDRAC device Object
- Create a Privilege Object
- Create an Association Object
- Add objects to an Association Object

Creating iDRAC device object

To create iDRAC device object:

- 1. In the MMC Console Root window, right-click a container.
- Select New > Dell Remote Management Object Advanced. The New Object window is displayed.
- **3.** Enter a name for the new object. The name must be identical to iDRAC name that you enter while configuring Active Directory properties using iDRAC Web interface.
- 4. Select iDRAC Device Object and click OK.

Creating privilege object

To create a privilege object:

(i) NOTE: You must create a privilege object in the same domain as the related association object.

- 1. In the Console Root (MMC) window, right-click a container.
- Select New > Dell Remote Management Object Advanced. The New Object window is displayed.
- 3. Enter a name for the new object.
- 4. Select Privilege Object and click OK.
- 5. Right-click the privilege object that you created, and select Properties.
- 6. Click the Remote Management Privileges tab and assign the privileges for the user or group.

Creating association object

To create association object: **(i) NOTE:** iDRAC association object is derived from the group and its scope is set to Domain Local.

1. In the Console Root (MMC) window, right-click a container.

- Select New > Dell Remote Management Object Advanced. This New Object window is displayed.
- 3. Enter a name for the new object and select Association Object.
- 4. Select the scope for the Association Object and click OK.
- 5. Provide access privileges to the authenticated users for accessing the created association objects.

158 Configuring user accounts and privileges



Providing user access privileges for association objects

To provide access privileges to the authenticated users for accessing the created association objects:

- 1. Go to Administrative Tools > ADSI Edit. The ADSI Edit window is displayed.
- 2. In the right-pane, navigate to the created association object, right-click and select **Properties**.
- 3. In the Security tab, click Add.
- 4. Type Authenticated Users, click Check Names, and click OK. The authenticated users is added to the list of Groups and user names.
- 5. Click OK.

Adding objects to association object

Using the **Association Object Properties** window, you can associate users or user groups, privilege objects, and iDRAC devices or iDRAC device groups.

You can add groups of users and iDRAC devices.

Adding users or user groups

To add users or user groups:

- 1. Right-click the Association Object and select Properties.
- 2. Select the Users tab and click Add.
- 3. Enter the user or user group name and click OK.

Adding privileges

To add privileges:

Click the **Privilege Object** tab to add the privilege object to the association that defines the user's or user group's privileges when authenticating to an iDRAC device. Only one privilege object can be added to an Association Object.

- 1. Select the Privileges Object tab and click Add.
- 2. Enter the privilege object name and click OK.
- 3. Click the **Privilege Object** tab to add the privilege object to the association that defines the user's or user group's privileges when authenticating to an iDRAC device. Only one privilege object can be added to an Association Object.

Adding iDRAC devices or iDRAC device groups

To add iDRAC devices or iDRAC device groups:

- 1. Select the **Products** tab and click **Add**.
- 2. Enter iDRAC devices or iDRAC device group name and click OK.
- 3. In the Properties window, click Apply and click OK.
- 4. Click the **Products** tab to add one iDRAC device connected to the network that is available for the defined users or user groups. You can add multiple iDRAC devices to an Association Object.

Configuring Active Directory with Extended schema using iDRAC web interface

To configure Active Directory with extended schema using Web interface:

(i) NOTE: For information about the various fields, see the *iDRAC Online Help*.

- In the iDRAC Web interface, go to iDRAC Settings > Users > Directory Services > Microsoft Active Directory. Click Edit
 - The Active Directory Configuration and Management Step 1 of 4 page is displayed.
- 2. Optionally, enable certificate validation and upload the CA-signed digital certificate used during initiation of SSL connections when communicating with the Active Directory (AD) server.



3. Click Next.

- The Active Directory Configuration and Management Step 2 of 4 page is displayed.
- Specify the location information about Active Directory (AD) servers and user accounts. Also, specify the time iDRAC must
 wait for responses from AD during login process.

```
() NOTE:
```

- If certificate validation is enabled, specify the Domain Controller Server addresses and the FQDN. Make sure that DNS is configured correctly under iDRAC Settings > Network
- If the user and iDRAC objects are in different domains, then do not select the **User Domain from Login** option. Instead select **Specify a Domain** option and enter the domain name where the iDRAC object is available.
- 5. Click Next. The Active Directory Configuration and Management Step 3 of 4 page is displayed.
- 6. Select Extended Schema and click Next. The Active Directory Configuration and Management Step 4 of 4 page is displayed.
- 7. Enter the name and location of the iDRAC device object in Active Directory (AD) and click **Finish**. The Active Directory settings for extended schema mode is configured.

Configuring Active Directory with Extended schema using RACADM

To configure Active Directory with Extended Schema using the RACADM:

1. Use the following commands:

```
racadm set iDRAC.ActiveDirectory.Enable 1
racadm set iDRAC.ActiveDirectory.Schema 2
racadm set iDRAC.ActiveDirectory.RacName <RAC common name>
racadm set iDRAC.ActiveDirectory.RacDomain <fully qualified rac domain name>
racadm set iDRAC.ActiveDirectory.DomainController1 <fully qualified domain name or IP
address of the domain controller>
racadm set iDRAC.ActiveDirectory.DomainController2 <fully qualified domain name or IP
address of the domain controller>
racadm set iDRAC.ActiveDirectory.DomainController3 <fully qualified domain name or IP
address of the domain controller>
```

- Enter the Fully Qualified Domain Name (FQDN) of the domain controller, not the FQDN of the domain. For example, enter servername.dell.com instead of dell.com.
- You must provide at least one of the three addresses. iDRAC attempts to connect to each of the configured addresses one-by-one until it makes a successful connection. With Extended Schema, these are the FQDN or IP addresses of the domain controllers where this iDRAC device is located.
- To disable the certificate validation during SSL handshake, use the following command:

```
racadm set iDRAC.ActiveDirectory.CertValidationEnable 0
```

In this case, you do not have to upload a CA certificate.

To enforce the certificate validation during SSL handshake (optional):

racadm set iDRAC.ActiveDirectory.CertValidationEnable 1

In this case, you must upload a CA certificate using the following command:

racadm sslcertupload -t 0x2 -f <ADS root CA certificate>

NOTE: If certificate validation is enabled, specify the Domain Controller Server addresses and the FQDN. Ensure that DNS is configured correctly under **iDRAC Settings** > **Network**.

Using the following RACADM command may be optional:

racadm sslcertdownload -t 1 -f <RAC SSL certificate>

2. If DHCP is enabled on iDRAC and you want to use the DNS provided by the DHCP server, enter the following command:

racadm set iDRAC.IPv4.DNSFromDHCP 1



3. If DHCP is disabled in iDRAC or you want to manually input your DNS IP address, enter the following command:

```
racadm set iDRAC.IPv4.DNSFromDHCP 0
racadm set iDRAC.IPv4.DNSFromDHCP.DNS1 <primary DNS IP address>
racadm set iDRAC.IPv4.DNSFromDHCP.DNS2 <secondary DNS IP address>
```

4. If you want to configure a list of user domains so that you only need to enter the user name during log in to iDRAC web interface, use the following command:

```
racadm set iDRAC.UserDomain.<index>.Name <fully qualified domain name or IP Address
of the domain controller>
```

You can configure up to 40 user domains with index numbers between 1 and 40.

Testing Active Directory settings

You can test the Active Directory settings to verify whether your configuration is correct, or to diagnose the problem with a failed Active Directory log in.

Testing Active Directory settings using iDRAC web interface

To test the Active Directory settings:

- In iDRAC Web Interface, go to iDRAC Settings > Users > Directory Services > Microsoft Active Directory, click Test. The Test Active Directory Settings page is displayed.
- 2. Click Test.
- Enter a test user's name (for example, username@domain.com) and password and click Start Test. A detailed test results and the test log displays.

If there is a failure in any step, examine the details in the test log to identify the problem and a possible solution.

NOTE: When testing Active Directory settings with Enable Certificate Validation checked, iDRAC requires that the Active Directory server be identified by the FQDN and not an IP address. If the Active Directory server is identified by an IP address, certificate validation fails because iDRAC is not able to communicate with the Active Directory server.

Testing Active Directory settings using RACADM

To test the Active Directory settings, use the testfeature command.

For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Configuring generic LDAP users

iDRAC provides a generic solution to support Lightweight Directory Access Protocol (LDAP)-based authentication. This feature does not require any schema extension on your directory services.

To make iDRAC LDAP implementation generic, the commonality between different directory services is utilized to group users and then map the user-group relationship. The directory service specific action is the schema. For example, they may have different attribute names for the group, user, and the link between the user and the group. These actions can be configured in iDRAC.

NOTE: The Smart Card based Two Factor Authentication (TFA) and the Single Sign-On (SSO) logins are not supported for generic LDAP Directory Service.



Configuring generic LDAP directory service using iDRAC webbased interface

To configure the generic LDAP directory service using Web interface:

(i) NOTE: For information about the various fields, see the *iDRAC Online Help*.

- In the iDRAC Web interface, go to iDRAC Settings > Users > Directory Services > Generic LDAP Directory Service, click Edit. The Generic LDAP Configuration and Management Step 1 of 3 page displays the current generic LDAP settings.
- Optionally, enable certificate validation and upload the digital certificate used during initiation of SSL connections when communicating with a generic LDAP server.

(i) NOTE: In this release, non-SSL port based LDAP bind is not supported. Only LDAP over SSL is supported.

3. Click Next.

- The Generic LDAP Configuration and Management Step 2 of 3 page is displayed.
- 4. Enable generic LDAP authentication and specify the location information about generic LDAP servers and user accounts.
 - (i) **NOTE:** If certificate validation is enabled, specify the LDAP Server's FQDN and make sure that DNS is configured correctly under **iDRAC Settings** > **Network**.
 - (i) **NOTE:** In this release, nested group is not supported. The firmware searches for the direct member of the group to match the user DN. Also, only single domain is supported. Cross domain is not supported.
- 5. Click Next.
 - The Generic LDAP Configuration and Management Step 3a of 3 page is displayed.
- 6. Click Role Group.
- The Generic LDAP Configuration and Management Step 3b of 3 page is displayed.
- 7. Specify the group distinguished name, the privileges associated with the group, and click Apply.
 - (i) **NOTE:** If you are using Novell eDirectory and if you have used these characters—#(hash), "(double quotes), ;(semi colon), > (greater than), , (comma), or <(lesser than)—for the Group DN name, they must be escaped.

The role group settings are saved. The **Generic LDAP Configuration and Management Step 3a of 3** page displays the role group settings.

- 8. If you want to configure additional role groups, repeat steps 7 and 8.
- 9. Click Finish. The generic LDAP directory service is configured.

Configuring generic LDAP directory service using RACADM

To configure the LDAP directory service, use the objects in the iDRAC.LDAP and iDRAC.LDAPRole groups.

For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Testing LDAP directory service settings

You can test the LDAP directory service settings to verify whether your configuration is correct, or to diagnose the problem with a failed LDAP log in.

Testing LDAP directory service settings using iDRAC web interface

To test the LDAP directory service settings:

- 1. In iDRAC Web Interface, go to **iDRAC Settings** > **Users** > **Directory Services** > **Generic LDAP Directory Service**. The **Generic LDAP Configuration and Management** page displays the current generic LDAP settings.
- 2. Click Test.
- **3.** Enter the user name and password of a directory user that is chosen to test the LDAP settings. The format depends on the *Attribute of User Login* is used and the user name entered must match the value of the chosen attribute.



- **NOTE:** When testing LDAP settings with **Enable Certificate Validation** checked, iDRAC requires that the LDAP server be identified by the FQDN and not an IP address. If the LDAP server is identified by an IP address, certificate validation fails because iDRAC is not able to communicate with the LDAP server.
- **NOTE:** When generic LDAP is enabled, iDRAC first tries to login the user as a directory user. If it fails, local user lookup is enabled.

The test results and the test log are displayed.

Testing LDAP directory service settings using RACADM

To test the LDAP directory service settings, use the testfeature command. For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.



System Configuration Lockdown mode

System Configuration Lockdown mode helps in preventing unintended changes after a system is provisioned. Lockdown mode is applicable to both configuration and firmware updates. When the system is locked down, any attempt to change the system configuration is blocked. If any attempts are made to change the critical system settings, an error message is displayed. Enabling System lockdown mode blocks the firmware update of third party I/O cards using the vendor tools.

System Lockdown mode is only available for Enterprise licensed customers.

In 4.40.00.00 release, System lockdown functionality is extended to NIC's also.

- (i) NOTE: Enhanced Lockdown for NIC's only includes firmware lockdown to prevent firmware updates. Configuration (x-UEFI) lockdown is not supported.
- () NOTE: After the System Lockdown mode is enabled, you cannot change any configuration settings. System settings fields are disabled.

Lockdown mode can be enabled or disabled using the following interfaces:

- iDRAC web interface
- RACADM
- WSMan
- SCP (System Configuration Profile)
- Redfish
- Using F2 during POST and selecting iDRAC Settings
- Factory System Erase

NOTE: To enable Lockdown mode, you must have iDRAC Enterprise or Datacenter license and Control and Configure system privileges.

NOTE: You may be able to access vMedia while system is in Lockdown mode but configuring remote file share is not enabled.

(i) NOTE: The interfaces like OMSA, SysCfg, and USC can only check the settings but cannot modify the configurations.

The following table lists the functional and nonfunctional features, interfaces, and utilities that are affected by Lockdown mode: **NOTE:** Changing the boot order using iDRAC is not supported when Lockdown mode is enabled. However, boot-control option is available in vConsole menu, which has no effect when iDRAC is in Lockdown mode.

Table 32. Items affected by Lockdown mode

Disabled	Remains functional
Deleting Licenses	Power Operations - Power ON/OFF, Reset Dewar con potition
SCP import	Power cap settingPower priority
Reset to defaults	Identify devices (Chassis or PERC)
OMSA/OMSS	• Part replacement, Easy Restore, and system board replacement
• IPMI	Running diagnostics
• DRAC/LC	Modular operations (FlexAddress or Remote-Assigned Address)
• DTK-Syscfg	Group Manager passcode
Redfish	• All vendor tools that have direct access to the device (excludes
OpenManage Essentials	selected NIC's)
 BIOS (F2 settings become read-only) 	License export
Group manager	• PERC
Select network cards	• PERC CLI
	• DTK-RAIDCFG
	◦ F2/Ctrl+R


Table 32. Items affected by Lockdown mode

Disabled	Remains functional
	 All Vendor tools that have direct access to the device NVMe DTK-RAIDCFG F2/Ctrl+R BOSS-S1 Marvell CLI F2/Ctrl+R ISM/OMSA settings (OS BMC enable, watchdog ping, OS name, OS version)

(i) NOTE: When lockdown mode is enabled, OpenID Connect login option is not displayed in iDRAC login page.

System Configuration Lockdown mode 165



Configuring iDRAC for Single Sign-On or smart card login

This section provides information to configure iDRAC for Smart Card login (for local users and Active Directory users), and Single Sign-On (SSO) login (for Active Directory users.) SSO and smart card login are licensed features.

iDRAC supports Kerberos based Active Directory authentication to support Smart Card and SSO logins. For information on Kerberos, see the Microsoft website.

Topics:

- Prerequisites for Active Directory Single Sign-On or smart card login
- Configuring iDRAC SSO login for Active Directory users
- Enabling or disabling smart card login
- Configuring Smart Card Login
- Using Smart Card to Login

Prerequisites for Active Directory Single Sign-On or smart card login

The prerequisites to Active Directory based SSO or Smart Card logins are:

- Synchronize iDRAC time with the Active Directory domain controller time. If not, kerberos authentication on iDRAC fails. You can use the Time zone and NTP feature to synchronize the time. To do this, see Configuring time zone and NTP.
- Register iDRAC as a computer in the Active Directory root domain.
- Generate a keytab file using the ktpass tool.
- To enable Single Sign-On for Extended schema, make sure that the Trust this user for delegation to any service (Kerberos only) option is selected on the Delegation tab for the keytab user. This tab is available only after creating the keytab file using ktpass utility.
- Configure the browser to enable SSO login.
- Create the Active Directory objects and provide the required privileges.
- For SSO, configure the reverse lookup zone on the DNS servers for the subnet where iDRAC resides.
 - (i) NOTE: If the host name does not match the reverse DNS lookup, Kerberos authentication fails.
- Configure the browser to support SSO login. For more information, see Single Sign-On.

(i) NOTE: Google Chrome and Safari do not support Active Directory for SSO login.

Registering iDRAC on Domain name System

To register iDRAC in Active Directory root domain:

- 1. Click **iDRAC Settings** > **Connectivity** > **Network**. The **Network** page is displayed.
- 2. You can select IPv4 Settings or IPv6 Settings based on the IP settings.
- **3.** Provide a valid **Preferred/Alternate DNS Server** IP address. This value is a valid DNS server IP address that is part of the root domain.
- 4. Select Register iDRAC on DNS.
- 5. Provide a valid DNS Domain Name.
- **6.** Verify that network DNS configuration matches with the Active Directory DNS information. For more information about the options, see the *iDRAC Online Help*.



Creating Active Directory objects and providing privileges

Logging in to Active Directory Standard schema based SSO

Perform the following steps for Active Directory Standard schema based SSO login:

- **1.** Create a User Group.
- **2.** Create a User for Standard schema.
- **NOTE:** Use the existing AD User Group & AD User.

Logging in to Active Directory Extended schema based SSO

Perform the following steps for Active Directory Extended schema based SSO login:

- 1. Create the device object, privilege object, and association object in the Active Directory server.
- 2. Set access privileges to the created privilege object.
 - (i) NOTE: It is recommended not to provide administrator privileges as this could bypass some security checks.
- **3.** Associate the device object and privilege object using the association object.
- 4. Add the preceding SSO user (login user) to the device object.
- 5. Provide access privilege to Authenticated Users for accessing the created association object.

Logging in to Active Directory SSO

Perform the following steps for Active Directory SSO login:

1. Create a Kerberos key-tab user which is used for the creation of the key-tab file.

(i) NOTE: Create new KERBROS key for every iDRAC IP.

Configuring iDRAC SSO login for Active Directory users

Before configuring iDRAC for Active Directory SSO login, make sure that you have completed all the prerequisites. You can configure iDRAC for Active Directory SSO when you setup an user account based on Active Directory.

Creating a User in Active Directory for SSO

To create a user in Active Directory for SSO:

- 1. Create a new user in the organization unit.
- 2. Go to Kerberos User>Properties>Account>Use Kerberos AES Encryption types for this account
- **3.** Use the following command to generate a Kerberos keytab in the Active Directory server:

```
C:\> ktpass.exe -princ HTTP/idrac7name.domainname.com@DOMAINNAME.COM -mapuser
DOMAINNAME\username -mapop set -crypto AES256-SHA1 -ptype KRB5_NT_PRINCIPAL -pass
[password] -out c:\krbkeytab
```

Note for Extended Schema

- Change the Delegation setting of the Kerberos user.
- Go to Kerberos User>Properties>Delegation>Trust this user for delegation to any service (Kerberos only)

(i) NOTE: Log-off and Log-in from the Management Station Active Directory user after changing the above setting.



Generating Kerberos keytab file

To support the SSO and smart card login authentication, iDRAC supports the configuration to enable itself as a kerberized service on a Windows Kerberos network. The Kerberos configuration on iDRAC involves the same steps as configuring a non–Windows Server Kerberos service as a security principal in Windows Server Active Directory.

The *ktpass* tool (available from Microsoft as part of the server installation CD/DVD) is used to create the Service Principal Name (SPN) bindings to a user account and export the trust information into a MIT–style Kerberos *keytab* file, which enables a trust relation between an external user or system and the Key Distribution Centre (KDC). The keytab file contains a cryptographic key, which is used to encrypt the information between the server and the KDC. The ktpass tool allows UNIX–based services that support Kerberos authentication to use the interoperability features provided by a Windows Server Kerberos KDC service. For more information on the **ktpass** utility, see the Microsoft website at: **technet.microsoft.com/en-us/library/cc779157(WS.10).aspx**

Before generating a keytab file, you must create an Active Directory user account for use with the **-mapuser** option of the ktpass command. Also, you must have the same name as iDRAC DNS name to which you upload the generated keytab file.

To generate a keytab file using the ktpass tool:

- 1. Run the *ktpass* utility on the domain controller (Active Directory server) where you want to map iDRAC to a user account in Active Directory.
- 2. Use the following ktpass command to create the Kerberos keytab file:

```
C:\> ktpass.exe -princ HTTP/idrac7name.domainname.com@DOMAINNAME.COM -mapuser
DOMAINNAME\username -mapop set -crypto AES256-SHA1 -ptype KRB5_NT_PRINCIPAL -pass
[password] -out c:\krbkeytab
```

The encryption type is AES256-SHA1. The principal type is KRB5_NT_PRINCIPAL. The properties of the user account to which the Service Principal Name is mapped to must have **Use AES 256 encryption types for this account** property enabled.

NOTE: Use lowercase letters for the **iDRACname** and **Service Principal Name**. Use uppercase letters for the domain name as shown in the example.

A keytab file is generated.

NOTE: If you find any issues with iDRAC user for which the keytab file is created, create a new user and a new keytab file. If the same keytab file which was initially created is again executed, it does not configure correctly.

Configuring iDRAC SSO login for Active Directory users using web interface

To configure iDRAC for Active Directory SSO login:

(i) NOTE: For information about the options, see the *iDRAC Online Help*.

- 1. Verify whether the iDRAC DNS name matches the iDRAC Fully Qualified Domain Name. To do this, in iDRAC Web interface, go to iDRAC Settings > Network > Common Settings and refer to DNS iDRAC Name property.
- 2. While configuring Active Directory to setup a user account based on standard schema or extended schema, perform the following two additional steps to configure SSO:
 - Upload the keytab file on the Active Directory Configuration and Management Step 1 of 4 page.
 - Select Enable Single Sign-On option on the Active Directory Configuration and Management Step 2 of 4 page.

Configuring iDRAC SSO login for Active Directory users using RACADM

To enable SSO, complete the steps to configure Active Directory, and run the following command:

racadm set iDRAC.ActiveDirectory.SSOEnable 1



Management Station Settings

Perform the following steps after configuring SSO login for Active Directory users:

- 1. Set the DNS Server IP in Network properties and mention the preferred DNS Server IP.
- 2. Go to My Computer and add the *domain.tld domain.
- 3. Add the Active Directory User to Administrator by navigating to: My Computer > Manage > Local User and Groups >
- **Groups** > **Administrator** and add the Active Directory User.
- Logoff the system and login using the Active Directory User credential.
 In Internet Explorer Setting, add *domain.tld domain as below:
 - a. Go to Tools > Internet Options > Security > Local Internet > Sites and clear the Automatically detect intranet network setting selection. Select the remaining three options and click Advanced to add *domain.tld.
 - b. Open a new window in IE and use the iDRAC hostname to launch the iDRAC GUI.
- 6. In Mozilla Firefox Setting, add *domain.tld domain:
 - Launch Firefox browser and type about:config in the URL.
 - Use negotiate in the filter textbox. Double click the result consisting of *auth.trusted.uris*. Type the domain, save the settings and close the browser.
 - Open a new window in Firefox and use the iDRAC hostname to launch the iDRAC GUI.

Enabling or disabling smart card login

Before enabling or disabling smart card login for iDRAC, make sure that:

- You have configure iDRAC permissions.
- iDRAC local user configuration or Active Directory user configuration with the appropriate certificates is complete.

NOTE: If smart card login is enabled, then SSH, IPMI Over LAN, Serial Over LAN, and remote RACADM are disabled. Again, if you disable smart card login, the interfaces are not enabled automatically.

Enabling or disabling smart card login using web interface

To enable or disable the Smart Card logon feature:

- In the iDRAC web interface, go to iDRAC Settings > Users > Smart Card. The Smart Card page is displayed.
- 2. From the **Configure Smart Card Logon** drop-down menu, select **Enabled** to enable smart card logon or select **Enabled With Remote RACADM**. Else, select **Disabled**.

For more information about the options, see the iDRAC Online Help.

 Click Apply to apply the settings. You are prompted for a Smart Card login during any subsequent logon attempts using the iDRAC web interface.

Enabling or disabling smart card login using RACADM

To enable smart card login, use the set command with objects in the iDRAC.SmartCard group.

For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Enabling or disabling smart card login using iDRAC settings utility

To enable or disable the Smart Card logon feature:

- 1. In the iDRAC Settings utility, go to Smart Card. The **iDRAC Settings Smart Card** page is displayed.
- 2. Select **Enabled** to enable smart card logon. Else, select **Disabled**. For more information about the options, see *iDRAC Settings Utility Online Help*.
- **3.** Click **Back**, click **Finish**, and then click **Yes**. The smart card logon feature is enabled or disabled based on the selection.



Configuring Smart Card Login

(i) **NOTE:** For Active Directory Smart Card Configuration, iDRAC must be configured either with Standard or Extended Schema SSO Login.

Configuring iDRAC smart card login for Active Directory users

Before configuring iDRAC Smart Card login for Active Directory users, make sure that you have completed the required prerequisites.

To configure iDRAC for smart card login:

- 1. In iDRAC Web interface, while configuring Active Directory to set up an user account based on standard schema or extended schema, on the Active Directory Configuration and Management Step 1 of 4 page:
 - Enable certificate validation.
 - Upload a trusted CA-signed certificate.
 - Upload the keytab file.
- 2. Enable smart card login. For information about the options, see the *iDRAC Online Help*.

Configuring iDRAC smart card login for local users

To configure iDRAC local user for smart card login:

- 1. Upload the smart card user certificate and trusted CA certificate to iDRAC.
- 2. Enable smart card login.

Uploading smart card user certificate

Before you upload the user certificate, make sure that the user certificate from the smart card vendor is exported in Base64 format. SHA-2 certificates are also supported.

Uploading smart card user certificate using web interface

To upload smart card user certificate:

1. In iDRAC web interface, go to iDRAC Settings > Users > Smart Card.

(i) NOTE: The Smart Card login feature requires the configuration of the local and/or Active Directory user certificate.

- 2. Under Configure Smart Card Logon, select Enabled With Remote RACADM to enable the configuration..
- 3. Set the option to Enable CRL Check for Smart Card Logon.
- 4. Click Apply.

Uploading smart card user certificate using RACADM

To upload smart card user certificate, use the **usercertupload** object. For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Requesting Certificate for smart card enrollment

Follow these steps to request certificate for smart card enrollment:

- 1. Connect the smart card in the client system and install the required drivers & software.
- 2. Verify the driver status in the Device Manager.
- 3. Launch the smart card enrollment agent in the browser.
- 4. Enter the Username & Password and click OK.
- 5. Click Request Certificate.



6. Click Advanced Certificate Request.

- 7. Click **Request a certificate** for a smart card on behalf of another user by using the smart card certificate enrollment station.
- 8. Select user to enroll by clicking **Select User** button.
- 9. Click **Enroll** and enter the smart card credential.
- 10. Enter the smart card PIN and click on **Submit**.

Uploading trusted CA certificate for smart card

Before you upload the CA certificate, make sure that you have a CA-signed certificate.

Uploading trusted CA certificate for smart card using web interface

To upload trusted CA certificate for smart card login:

- In iDRAC Web interface, go to iDRAC Settings > Network > User Authentication > Local Users. The Users page is displayed.
- 2. In the User ID column, click a user ID number. The Users Main Menu page is displayed.
- **3.** Under Smart Card Configurations, select Upload Trusted CA Certificate and click Next. The Trusted CA Certificate Upload page is displayed.
- 4. Browse and select the trusted CA certificate, and click **Apply**.

Uploading trusted CA certificate for smart card using RACADM

To upload trusted CA certificate for smart card login, use the **usercertupload** object. For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Using Smart Card to Login

(i) NOTE: Smart card login is supported only on Internet Explorer.

To login using smart card:

- 1. Logout from iDRAC GUI after enabling smart card.
- 2. Launch iDRAC by using http://IP/ or launch using FQDN http://FQDN/
- 3. Click Install after smart card plug-in download.
- 4. Enter smart card PIN and click Submit.
- **5.** iDRAC will login successfully using smart card.



Configuring iDRAC to send alerts

You can set alerts and actions for certain events that occur on the managed system. An event occurs when the status of a system component is greater than the pre-defined condition. If an event matches an event filter and you have configured this filter to generate an alert (e-mail, SNMP trap, IPMI alert, remote system logs, Redfish event, or WS events), then an alert is sent to one or more configured destinations. If the same event filter is also configured to perform an action (such as reboot, power cycle, or power off the system), the action is performed. You can set only one action for each event.

To configure iDRAC to send alerts:

- 1. Enable alerts.
- 2. Optionally, you can filter the alerts based on category or severity.
- **3.** Configure the e-mail alert, IPMI alert, SNMP trap, remote system log, Redfish event, operating system log, and/or WS-event settings.
- 4. Enable event alerts and actions such as:
 - Send an email alert, IPMI alert, SNMP traps, remote system logs, Redfish event, operating system log, or WS events to configured destinations.
 - Perform a reboot, power off, or power cycle the managed system.

Topics:

- Enabling or disabling alerts
- Filtering alerts
- Setting event alerts
- Setting alert recurrence event
- Setting event actions
- Configuring email alert, SNMP trap, or IPMI trap settings
- Configuring WS Eventing
- Configuring Redfish Eventing
- Monitoring chassis events
- Alerts message IDs

Enabling or disabling alerts

For sending an alert to configured destinations or to perform an event action, you must enable the global alerting option. This property overrides individual alerting or event actions that is set.

Enabling or disabling alerts using web interface

To enable or disable generating alerts:

- In iDRAC web interface, go to Configuration > System Settings > Alert Configuration. The Alerts page is displayed.
- 2. Under Alerts section:
 - Select Enable to enable alert generation or perform an event action.
 - Select **Disable** to disable alert generation or disable an event action.
- 3. Click Apply to save the setting.

Quick Alert Configuration

To configure alerts in bulk:

1. Go to Quick Alert Configuration under Alert Configuration page.

172 Configuring iDRAC to send alerts



2. Under Quick Alert Configuration section:

- Select the alert category.
- Select the issue severity notification.
- Select the location where you would like to receive these notifications.
- 3. Click Apply to save the setting.

(i) NOTE: You must select at least one category, one severity, and one destination type to apply the configuration.

All the alerts that are configured are displayed in total under **Alerts Configuration Summary**.

Enabling or disabling alerts using RACADM

Use the following command:

racadm set iDRAC.IPMILan.AlertEnable <n>

```
n=0 — Disabled
```

n=1 — Enabled

Enabling or disabling alerts using iDRAC settings utility

To enable or disable generating alerts or event actions:

- In the iDRAC Settings utility, go to Alerts. The iDRAC Settings Alerts page is displayed.
- 2. Under **Platform Events**, select **Enabled** to enable alert generation or event action. Else, select **Disabled**. For more information about the options, see *iDRAC Settings Utility Online Help*.
- **3.** Click **Back**, click **Finish**, and then click **Yes**. The alert settings are configured.

Filtering alerts

You can filter alerts based on category and severity.

Filtering alerts using iDRAC web interface

To filter the alerts based on category and severity:

(i) NOTE: Even if you are a user with read-only privileges, you can filter the alerts.

- 1. In iDRAC Web interface, go to Configuration > System Settings > Alerts and Remote System Log Configuration.
- 2. Under Alerts and Remote System Log Configuration section, select Filter:
 - System Health System Health category represents all the alerts that are related to hardware within the system chassis. Examples include temperature errors, voltage errors, device errors.
 - Storage Health Storage Health category represents alerts that are related to the storage subsystem. Examples include, controller errors, physical disk errors, virtual disk errors.
 - Configuration Configuration category represents alerts that are related to hardware, firmware and software configuration changes. Examples include, PCI-e card added/removed, RAID configuration changed, iDRAC license changed.
 - Audit Audit category represents the audit log. Examples include, user login/logout information, Password authentication failures, session info, power states.
 - Updates Update category represents alerts that are generated due to firmware/Driver upgrades/downgrades.
 NOTE: This doesn't represent firmware inventory.
 - Work Notes
- 3. Select one or more of the following severity levels:

Configuring iDRAC to send alerts 173



- Informational
- Warning
- Critical
- 4. Click Apply.
 - The **Alert Results** section displays the results based on the selected category and severity.

Filtering alerts using RACADM

To filter the alerts, use the **eventfilters** command. For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Setting event alerts

You can set event alerts such as e-mail alerts, IPMI alerts, SNMP traps, remote system logs, operating system logs, and WS events to be sent to configured destinations.

Setting event alerts using web interface

To set an event alert using the web interface:

- 1. Make sure that you have configured the e-mail alert, IPMI alert, SNMP trap settings, and/or remote system log settings.
- 2. In iDRAC Web interface, go to Configuration > System Settings > Alerts and Remote System Log Configuration.
- 3. Under Category, select one or all of the following alerts for the required events:
 - Email
 - SNMP Trap
 - IPMI Alert
 - Remote System Log
 - WS Eventing
 - OS Log
 - Redfish Event
- 4. Select Action.
 - The setting is saved.
- 5. Optionally, you can send a test event. In the **Message ID to Test Event** field, enter the message ID to test if the alert is generated and click **Test**. For more information about the event and error messages generated by the system firmware and agents that monitor system components, see the *Event and Error Message Reference Guide* at iDRACmanuals

Setting event alerts using RACADM

To set an event alert, use the **eventfilters** command. For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Setting alert recurrence event

You can configure iDRAC to generate additional events at specific intervals if the system continues to operate at a temperature which is greater than the inlet temperature threshold limit. The default interval is 30 days. The valid range is 0 to 366 days. A value of '0' indicates no event recurrence.

(i) NOTE: You must have Configure iDRAC privilege to set the alert recurrence value.

Setting alert recurrence events using RACADM

To set the alert recurrence event using RACADM, use the **eventfilters** command. For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

174 Configuring iDRAC to send alerts

Inserido ao protocolo **22.951.206-4** por: **Pedro Henrique Golin Linhares** em: 31/03/2025 12:49. A autenticidade deste documento pode ser validada no endereço: https://www.eprotocolo.pr.gov.br/spiweb/validarDocumento com o código: fc1c44dfe0686d5547ffd691199be699.



Setting alert recurrence events using iDRAC web interface

To set the alert recurrence value:

- 1. In iDRAC Web interface, go to Configuration > System Settings > Alert Recurrence.
- 2. In the **Recurrence** column, enter the alert frequency value for the required category, alert, and severity type(s). For more information, see the *iDRAC Online help*.
- 3. Click Apply.

The alert recurrence settings are saved.

Setting event actions

You can set event actions such as perform a reboot, power cycle, power off, or perform no action on the system.

Setting event actions using web interface

To set an event action:

- 1. In iDRAC Web interface, go to Configuration > System Settings > Alert and Remote System Log Configuration.
- 2. From the Actions drop-down menu, for each event select an action:
 - Reboot
 - Power Cycle
 - Power Off
 - No Action
- Click Apply. The setting is saved.

Setting event actions using RACADM

To configure an event action, use the eventfilters command. For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Configuring email alert, SNMP trap, or IPMI trap settings

The management station uses Simple Network Management Protocol (SNMP) and Intelligent Platform Management Interface (IPMI) traps to receive data from iDRAC. For systems with large number of nodes, it may not be efficient for a management station to poll each iDRAC for every condition that may occur. For example, event traps can help a management station with load balancing between nodes or by issuing an alert if an authentication failure occurs. SNMP v1, v2, and v3 formats are supported.

You can configure the IPv4 and IPv6 alert destinations, email settings, and SMTP server settings, and test these settings. You can also specify the SNMP v3 user to whom you want to send the SNMP traps.

Before configuring the email, SNMP, or IPMI trap settings, make sure that:

- You have Configure RAC permission.
- You have configured the event filters.

Configuring IP alert destinations

You can configure the IPv6 or IPv4 addresses to receive the IPMI alerts or SNMP traps.

For information about the iDRAC MIBs required to monitor the servers using SNMP, see the *Dell EMC OpenManage SNMP Reference Guide* available at https://www.dell.com/openmanagemanuals.

Configuring iDRAC to send alerts 175



Configuring IP alert destinations using web interface

To configure alert destination settings using Web interface:

- 1. In iDRAC Web interface, go to Configuration > System Settings > SNMP and E-mail Settings.
- Select the State option to enable an alert destination (IPv4 address, IPv6 address, or Fully Qualified Domain Name (FQDN)) to receive the traps.

You can specify up to eight destination addresses. For more information about the options, see the *iDRAC Online Help*.

- **3.** Select the SNMP v3 user to whom you want to send the SNMP trap.
- **4.** Enter the iDRAC SNMP community string (applicable only for SNMPv1 and v2) and the SNMP alert port number. For more information about the options, see the *iDRAC Online Help*.
 - () NOTE: The Community String value indicates the community string to use in a Simple Network Management Protocol (SNMP) alert trap sent from iDRAC. Make sure that the destination community string is the same as the iDRAC community string. The default value is Public.
- To test whether the IP address is receiving the IPMI or SNMP traps, click Send under Test IPMI Trap and Test SNMP Trap respectively.
- 6. Click Apply.
 - The alert destinations are configured.
- In the SNMP Trap Format section, select the protocol version to be used to send the traps on the trap destination(s) SNMP v1, SNMP v2, or SNMP v3 and click Apply.

NOTE: The **SNMP Trap Format** option applies only for SNMP Traps and not for IPMI Traps. IPMI Traps are always sent in SNMP v1 format and is not based on the configured **SNMP Trap Format** option.

The SNMP trap format is configured.

Configuring IP alert destinations using RACADM

To configure the trap alert settings:

1. To enable traps:

```
racadm set idrac.SNMP.Alert.<index>.Enable <n>
```

Parameter	Description	
<index></index>	Destination index. Allowed values are 1 through 8.	
<n>=0</n>	Disable the trap	
<n>=1</n>	Enable the trap	

2. To configure the trap destination address:

```
racadm set idrac.SNMP.Alert.<index>.DestAddr <Address>
```

Parameter Description	
<index></index>	Destination index. Allowed values are 1 through 8.
<address></address>	A valid IPv4, IPv6, or FQDN address

3. Configure the SNMP community name string:

racadm set idrac.ipmilan.communityname <community_name>

Parameter	Description
<community_name></community_name>	The SNMP Community Name.

4. To configure SNMP destination:

• Set the SNMP trap destination for SNMPv3:



racadm set idrac.SNMP.Alert.<index>.DestAddr <IP address>

• Set SNMPv3 users for trap destinations:

racadm set idrac.SNMP.Alert.<index>.SNMPv3Username <user name>

• Enable SNMPv3 for a user:

racadm set idrac.users.<index>.SNMPv3Enable Enabled

5. To test the trap, if required:

racadm testtrap -i <index>

For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Configuring IP alert destinations using iDRAC settings utility

You can configure alert destinations (IPv4, IPv6, or FQDN) using the iDRAC Settings utility. To do this:

- In the iDRAC Settings utility, go to Alerts. The iDRAC Settings Alerts page is displayed.
- 2. Under **Trap Settings**, enable the IP address(es) to receive the traps and enter the IPv4, IPv6, or FQDN destination address(es). You can specify up to eight addresses.
- **3.** Enter the community string name. For information about the options, see the *iDRAC Settings Utility Online Help*.
- **4.** Click **Back**, click **Finish**, and then click **Yes**. The alert destinations are configured.

Configuring email alert settings

You can configure the sender email address and receiver (destination) email address to receive the email alerts. Also, configure the SMTP server address settings.

- () NOTE: Email alerts support both IPv4 and IPv6 addresses. The iDRAC DNS Domain Name must be specified when using IPv6.
- **NOTE:** If you are using an external SMTP server, ensure that iDRAC can communicate with that server. If the server is unreachable, the error RAC0225 is displayed while trying to send a test mail.

Configuring email alert settings using web interface

To configure the email alert settings using Web interface:

- 1. In iDRAC Web interface, go to Configuration > System Settings > SMTP (E-mail) Configuration.
- **2.** Type a valid email address.
- 3. Click Send under Test Email to test the configured email alert settings.
- 4. Click Apply.
- 5. For SMTP (E-mail) Server Settings provide the following details:
 - SMTP (E-mail) Server IP Address or FQDN/DNS Name
 - Custom Sender Address This field has the following options:
 - **Default** Address field is not editable
 - \circ **Custom** You can enter the email ID from which you can receive the email alerts
 - Custom Message Subject Prefix This field has the following options:
 - **Default** Default message is not editable
 - $\circ~$ ${\rm Custom}-{\rm You}$ can choose the message to appear in the ${\rm Subject}$ line of the email
 - SMTP Port Number The connection can be encrypted and emails can be sent over secure ports:

Configuring iDRAC to send alerts 177



- No Encryption port 25 (default)
- SSL Port 465
- Connection Encryption When you do not have an email server in your premises, you can use cloud based email servers or SMTP Relays. To configure cloud email server, you can set this feature to any of the following values from the drop down:
 - None No encryption on the connection to the SMTP server. It is the default value.
 - **SSL** Runs SMTP protocol over SSL

() NOTE:

- This feature is not configurable via Group Manager.
- This is a licensed feature and is not available in iDRAC Basic License.
- You must have Configure iDARC privilege to use this feature.
- Authentication
- Username

For Server settings, the port usage depends on connectionencryptiontype and this can be configured only using RACADM.

6. Click Apply. For more information about the options, see the iDRAC Online Help.

Configuring email alert settings using RACADM

1. To enable email alert:

```
racadm set iDRAC.EmailAlert.Enable.[index] [n]
```

Parameter	Description	
index	Email destination index. Allowed values are 1 through 4.	
n=0	Disables email alerts.	
n=1	Enables email alerts.	

2. To configure email settings:

```
racadm set iDRAC.EmailAlert.Address.[index] [email-address]
```

Parameter	Description
index	Email destination index. Allowed values are 1 through 4.
email-address Destination email address that receives the platform event alerts.	

3. To configure sender email settings:

```
racadm set iDRAC.RemoteHosts.[index] [email-address]
```

Parameter	Description
index	Sender Email index.
email-address	Sender email address that sends the platform event alerts.

4. To configure a custom message:

```
racadm set iDRAC.EmailAlert.CustomMsg.[index] [custom-message]
```

Parameter	Description	
index	Email destination index. Allowed values are 1 through 4.	

/	POTOCOLO
	Fls.1221 Mov. 33
	PRADO DO ES

Parameter	Description	Ì
custom-message	Custom message	

5. To test the configured email alert, if required:

racadm testemail -i [index]

Parameter	Description	
index	Email destination index to be tested. Allowed values are 1 through 4.	

For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Configuring SMTP email server address settings

You must configure the SMTP server address for email alerts to be sent to specified destinations.

Configuring SMTP email server address settings using iDRAC web interface

To configure the SMTP server address:

- 1. In iDRAC Web interface, go to Configuration > System Settings > Alert Configuration > SNMP (E-mail Configuration).
- 2. Enter the valid IP address or fully qualified domain name (FQDN) of the SMTP server to be used in the configuration.
- **3.** Select the **Enable Authentication** option and then provide the user name and password (of a user who has access to SMTP server).
- Enter the SMTP port number.
 For more information about the fields, see the *iDRAC Online Help*.
- 5. Click **Apply**. The SMTP settings are configured.

Configuring SMTP email server address settings using RACADM

To configure the SMTP email server:

racadm set iDRAC.RemoteHosts.SMTPServerIPAddress <SMTP E-mail Server IP Address>

Configuring WS Eventing

The WS Eventing protocol is used for a client service (subscriber) to register interest (subscription) with a server (event source) for receiving messages containing the server events (notifications or event messages). Clients interested in receiving the WS Eventing messages can subscribe with iDRAC and receive Lifecycle Controller job related events.

The steps required to configure WS Eventing feature to receive WS Eventing messages for changes related to Lifecycle Controller jobs are described in the Web service Eventing Support for iDRAC 1.30.30 specification document. In addition to this specification, see the DSP0226 (DMTF WS Management Specification), Section 10 Notifications (Eventing) document for the complete information on the WS Eventing protocol. The Lifecycle Controller related jobs are described in the DCIM Job Control Profile document.

Configuring Redfish Eventing

The Redfish eventing protocol is used for a client service (subscriber) to register interest (subscription) with a server (event source) for receiving messages containing the Redfish events (notifications or event messages). Clients interested in receiving the Redfish eventing messages can subscribe with iDRAC and receive Lifecycle Controller job related events.



Monitoring chassis events

On the PowerEdge FX2/FX2s chassis, you can enable the **Chassis Management and Monitoring** setting in iDRAC to perform chassis management and monitoring tasks such as monitoring chassis components, configuring alerts, using iDRAC RACADM to pass CMC RACADM commands, and updating the chassis management firmware. This setting allows you to manage the servers in the chassis even if the CMC is not on the network. You can set the value to **Disabled** to forward the chassis events. By default, this setting is set as **Enabled**.

(i) NOTE: For this setting to take effect, you must ensure that in CMC, the Chassis Management at Server setting must be set to Monitor or Manage and Monitor.

When the **Chassis Management and Monitoring** option is set to **Enabled**, iDRAC generates and logs chassis events. The events generated are integrated into the iDRAC event subsystem and alerts are generated similar to the rest of the events.

CMC also forwards the events generated to iDRAC. In case the iDRAC on the server is not functional, CMC queues the first 16 events and logs the rest in the CMC log. These 16 events are sent to iDRAC as soon as **Chassis monitoring** is set to enabled.

In instances where iDRAC detects that a required CMC functionality is absent, a warning message is displayed informing you that certain features may not be functional without a CMC firmware upgrade.

(i) NOTE: iDRAC doesnot support the following Chassis attributes:

- ChassisBoardPartNumber
- ChassisBoardSerialNumber

Monitoring chassis events using the iDRAC web interface

To monitor chassis events using the iDRAC web interface, perform the following steps:

NOTE: This section appears only for PowerEdge FX2/FX2s chassis and if **Chassis Management at Server** mode is set to **Monitor** or **Manage and Monitor** in CMC.

- 1. On the CMC interface, click Chassis Overview > Setup > General.
- 2. From the Chassis Management at Server Mode drop-down menu, select Manage and Monitor, and click Apply.
- 3. Launch the iDRAC web interface, click **Overview** > iDRAC Settings > CMC.
- 4. Under the Chassis Management at Server section, ensure that Capability from iDRAC drop-down box is set to Enabled.

Monitoring chassis events using RACADM

This setting is applicable only for PowerEdge FX2/FX2s servers and if **Chassis Management at Server** mode is set to **Monitor** or **Manage and Monitor** in CMC.

To monitor chassis events using iDRAC RACADM:

racadm get system.chassiscontrol.chassismanagementmonitoring

For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Alerts message IDs

The following table provides the list of message IDs that are displayed for the alerts.

Table 33. Alert message IDs

Message ID	Description	Description (For MX platforms)
АМР	Amperage	Amperage
ASR	Auto Sys Reset	Auto Sys Reset



Table 33. Alert message IDs (continued)

Message ID	Description	Description (For MX platforms)
ВАТ	Battery Event	Battery Event
BIOS	BIOS Management	BIOS Management
BOOT	BOOT Control	BOOT Control
CBL	Cable	Cable
CPU	Processor	Processor
CPUA	Proc Absent	Proc Absent
CTL	Storage Contr	Storage Contr
DH	Cert Mgmt	Cert Mgmt
DIS	Auto-Discovery	Auto-Discovery
ENC	Storage Enclosr	Storage Enclosr
FAN	Fan Event	Fan Event
FSD	Debug	Debug
HWC	Hardware Config	Hardware Config
IPA	DRAC IP Change	DRAC IP Change
ITR	Intrusion	Intrusion
JCP	Job Control	Job Control
LC	Lifecycle Controller	Lifecycle Controller
LIC	Licensing	Licensing
LNK	Link Status	Link Status
LOG	Log event	Log event
MEM	Memory	Memory
NDR	NIC OS Driver	NIC OS Driver
NIC	NIC Config	NIC Config
OSD	OS Deployment	OS Deployment
OSE	OS Event	OS Event
PCI	PCI Device	PCI Device
PDR	Physical Disk	Physical Disk
PR	Part Exchange	Part Exchange
PST	BIOS POST	BIOS POST

Configuring iDRAC to send alerts 181



Table 33. Alert message IDs (continued)

Message ID	Description	Description (For MX platforms)
PSU	Power Supply	Power Supply
PSUA	PSU Absent	PSU Absent
PWR	Power Usage	Power Usage
RAC	RAC Event	RAC Event
RDU	Redundancy	Redundancy
RED	FW Download	FW Download
RFL	IDSDM Media	IDSDM Media
RFLA	IDSDM Absent	IDSDM Absent
RFM	FlexAddress SD	Not Applicable
RRDU	IDSDM Redundancy	IDSDM Redundancy
RSI	Remote Service	Remote Service
SEC	Security Event	Security Event
SEL	Sys Event Log	Sys Event Log
SRD	Software RAID	Software RAID
SSD	PCIe SSD	PCle SSD
STOR	Storage	Storage
SUP	FW Update Job	FW Update Job
SWC	Software Config	Software Config
SWU	Software Change	Software Change
SYS	System Info	System Info
ТМР	Temperature	Temperature
TST	Test Alert	Test Alert
UEFI	UEFI Event	UEFI Event
USR	User Tracking	User Tracking
VDR	Virtual Disk	Virtual Disk
VF	vFlash SD card	vFlash SD card
VFL	vFlash Event	vFlash Event
VFLA	vFlash Absent	vFlash Absent
VLT	Voltage	Voltage

182 Configuring iDRAC to send alerts

Inserido ao protocolo **22.951.206-4** por: **Pedro Henrique Golin Linhares** em: 31/03/2025 12:49. A autenticidade deste documento pode ser validada no endereço: https://www.eprotocolo.pr.gov.br/spiweb/validarDocumento com o código: fc1c44dfe0686d5547ffd691199be699.



Table 33. Alert message IDs (continued)

Message ID	Description	Description (For MX platforms)
VME	Virtual Media	Virtual Media
VRM	Virtual Console	Virtual Console
WRK	Work Note	Work Note



iDRAC 9 Group Manager

Group Manager enables user to have multiple console experience and offers simplified basic iDRAC management.

iDRAC Group Manager feature is available for Dell's 14th generation servers to offer simplified basic management of iDRACs and associated servers on the local network using the iDRAC GUI. Group Manager allows 1XMany console experience without involving a separate application. It allows the users to view the details of a set of servers by permitting more powerful management than by inspecting servers visually for faults and other manual methods.

Group manager is a licensed feature and part of the Enterprise license. Only iDRAC admin users can access the Group Manager functionality.

(i) NOTE: For better user experience Group Manager supports up to 250 server nodes.

Topics:

- Group Manager
- Summary View
- Network Configuration requirements
- Manage Logins
- Configure Alerts
- Export
- Discovered Servers View
- Jobs View
- Jobs Export
- Group Information Panel
- Group Settings
- Actions on a selected Server
- iDRAC Group Firmware Update

Group Manager

To use **Group Manager** feature, you need to enable the **Group Manager** from iDRAC index page or on the Group Manager Welcome screen. The group manager welcome screen provides options listed in the below table.

Table 34. Options in Group Manager

Option	Description
Join Existing Group	 Allows you to join an existing group, you need to know the GroupName and Passcode to join a specific group. (i) NOTE: Passwords are associated to iDRAC user credentials. Whereas, a passcode is associated to a group to establish authenticated device communication between different iDRACs in the same group.
Create New Group	Allows you to create a new group. The specific iDRAC which has created the group would be the master (primary controller) of the group.
Disable Group Manager for this System	You can select this option in case you do not want to join any group from a specific system. However, you can access Group Manager at any point of time by selecting Open Group Manager from the iDRAC index page. Once you disable the group manager, user needs to wait for 60 seconds to perform any further group manager operations.



Once the group manager feature is enabled, that iDRAC allows you the option to create or join an iDRAC local group. More than one iDRAC group can be setup in the local network but an individual iDRAC can only be a member of one group at a time. To change group (join a new group) the iDRAC must first leave its current group and then join the new group. The iDRAC from where the group was created gets chosen as the primary controller of the group by default. The user does not define a dedicated group manager primary controller to control that group. The primary controller hosts the group manager web interface and provides the GUI based work flows. The iDRAC members self-select a new primary controller for the group if the current primary goes offline for a prolonged duration, but that does not have any impact on the end user. You can normally access the group manager from all iDRAC members by clicking group manager from the iDRAC index page.

Summary View

You need to have administrator privileges to access group manager pages. If a non-administrator user logs onto the iDRAC, the group manager section does not appear with their credentials. The group manager home page (summary view) is broadly categorized as three sections. The first section shows rollup summary with aggregated summary details.

- Total number of servers in the local group.
- Chart showing number of servers per server model.
- Doughnut chart showing the servers per their health status (clicking on a chart section filters the server list to show only the servers with selected health).
- Warning box if there is a duplicate group detected in the local network. Duplicate group is typically the group with the same name but with a different passcode. This warning box does not appear if there is no duplicate group.
- Displays the iDRACs, that are controlling the group (Primary and Secondary controller).

The second section provides buttons for actions that are taken on the group as whole and the third section displays the list of all iDRACs in the group.

It shows all the systems in the group and their current health status and allows the user to take corrective action as needed. Server attributes specific to a server is described in the below table.

Table 35. Server Attributes

Server Attribute	Description
Health	Indicates the health status of that specific server.
Host Name	Displays the Server name.
iDRAC IP Address	Displays the exact IPV4 and IPV6 addresses.
Service Tag	Displays the Service Tag information.
Model	Displays the Model number of the Dell Server.
idrac	Displays the iDRAC version.
Last Status Update	Displays the time stamp when the server Status was last updated.

The System Information panel provides further details on the server like iDRAC network connectivity status, server host power state, express service code, operating system, asset tag, node ID, IDRAC DNS name, Server BIOS version, Server CPU information, System memory and location information. You may double click on a row or click on the launch iDRAC button to perform a single sign on redirect to the selected iDRAC index page. On the selected server, virtual console can be accessed or server power actions can be performed from More Actions dropped down list.

Manage iDRAC user logins, Alert configuration and group inventory export are the group actions supported.

Network Configuration requirements

Group Manager uses IPv6 link local networking to communicate between iDRACs (excluding the web browser GUI). Link local communication is defined as non-routed packets which means any iDRAC separated by a router cannot be joined in a local group. If the iDRAC-dedicated port or shared LOM is assigned to a vLAN, the vLAN limits the number of iDRACs that can be joined in a group (iDRACs must be on same vLAN and traffic must not pass through a router).

When Group Manager is enabled, iDRAC enables an IPv6 Link Local address regardless of the iDRAC's current user defined network configuration. Group Manager can be used when iDRAC is configured for IPv4 or IPv6 IP addresses.



Group Manager uses mDNS to discover other iDRACs on the network and sends encrypted packets for normal inventorying, monitoring and management of the group using the link local IP address. Using IPv6 link local networking means that the Group Manager ports and packets will never leave the local network or be accessible to external networks.

Ports (Specific to Group Manager unique functionality does not include all iDRAC ports) are:

- 5353 (mDNS)
- 443 (webserver) configurable
- 5670 (Multicast group communication)
- C000 -> F000 dynamically identifies one free port for each member to communicate in the group

Best networking practices

- Groups are intended to be small and on the same physical link local network.
- It is recommended to use the dedicated iDRAC network port for enhanced security. Shared LOM is also supported.

Additional network considerations

Two iDRACs that are separated by a router in the network topology are considered to be on separate local networks and cannot be joined in the same iDRAC local group. Meaning, if the iDRAC is configured for dedicated NIC settings, the network cable connected to iDRAC dedicated port in the rear of the server must be under a local network for all relevant servers.

If the iDRAC is configured for shared LOM network settings, the shared network connection used by both server host and IDRAC need to be connected under a local network for Group Manager to detect and onboard those servers into a common group. IDRACs configured with a mix of dedicated and shared LOM mode NIC settings could also be on-boarded into a common group, if all the network connections do not pass through a router.

Effect of MLD snooping in VLAN environments on Group Manager Discovery

Since Group Manager uses IPv6 multicast addressing for node-initiated discovery, a feature called MLD Snooping can prevent Group Manager-enabled devices from discovering each other if not configured properly. MLD Snooping is a common ether switch feature intended to reduce the amount of unnecessary IPv6 multicast traffic on a network.

If MLD Snooping is active in any network, ensure there is an MLD querier enabled so that the ether switches are kept up to date with the active Group Manager devices on the network. Alternatively, if MLD Snooping is not needed, it can be disabled. Note that some network switches have MLD Snooping enabled by default. And it is same for switching modules in the MX7000 chassis.

(i) NOTE:

For example

• To disable MLD snooping on a VLAN on a MX5108n IOM:

MX5108N-B1# configure terminal

MX5108N-B1(config)# interface vlan 194

MX5108N-B1(conf-if-vl-194)#no ipv6 mld snooping

• To enable an MLD querier in a VLAN on the MX5108n IOM:

MX5108N-B1# configure terminal

MX5108N-B1(config)# interface vlan 194

MX5108N-B1(conf-if-vl-194)#ipv6 mld snooping querier

Manage Logins

Use this section to Add New User, Change User Password and Delete User from the Group.



Group jobs including Manage Logins are one time configurations of the servers. Group manager uses SCP and jobs to make any changes. Every iDRAC in the group owns an individual job in its job queue for each Group Manager job. Group Manager does not detect changes on member iDRACs or lock member configurations.

(i) NOTE: Group jobs does not configure or override the lockdown mode for any specific iDRAC.

Leaving a group does not change local user or change settings on a member iDRAC.

Add a New User

Use this section to create and add a new user profile on all the servers in that group. A group job would be created to add the user to all servers in that group. The status of group job can be found at **GroupManager** > **Jobs** page.

(i) **NOTE:** By default iDRAC is configured with a local administrator account. You can access further information for each parameter with local administrator account.

For more information see, Configuring user accounts and privileges.

Table 36. New User Options

Option	Description
New User Information	Allows you to provide the new user's information details.
iDRAC Permissions	Allows you to define the user's role for future usage.
Advanced User Settings	Allows you to set (IPMI) user privileges and helps you to enable SNMP.

NOTE: Any member iDRAC with system lockdown enabled, that is part of the same group returns an error that the user password was not updated.

Change User Password

Use this section to change the password information for the user. You can see the **Users** detail with the **User Name**, **Role** and **Domain** information for individual user. A group job would be created to change the user password on all the servers in that group. The status of group job can be found at **GroupManager** > **Jobs** page.

If user already exists then the password can be updated. Any member iDRAC with system lockdown enabled, that is part of the group returns an error that the user password was not updated. If the user does not exist, then an error is returned to group manager indicating that the user does not exist on the system. The list of users shown in Group Manager GUI is based on the current user list on the iDRAC that is acting as the primary controller. It does not show all users for all iDRACs.

Delete User

Use this section to delete users from all the group servers. A group job would be created to delete users from all the group servers. The status of group job can be found at **GroupManager** > **Jobs** page.

If user already exists on a member iDRAC then the user can be deleted. Any member iDRAC with system lockdown enabled that is part of the group returns an error that the user is not deleted. If the user does not exist then it shows a successful deletion for that iDRAC. The list of users shown in Group Manager GUI is based on the current user list on the iDRAC which is acting as the primary controller. It does not show all users for all iDRACs.

Configure Alerts

Use this section to configure e-mail alerts. By default alerting is disabled. However, you can enable the alerting anytime. A group job would be created to apply the e-mail alert configuration to all the group servers. The status of group job can be monitored at **GroupManager** > **Jobs** page. Group manager email alert configures email alerts on all members. It sets the SMTP server settings on all members in the same group. Each iDRAC is configured separately. Email configuration is not globally saved. Current values are based on the iDRAC that is acting as a primary controller. Leaving a group does not reconfigure email alerts.

For more information on Configuring Alerts, see Configuring iDRAC to send alerts.



Table 37. Configuring alerts options

Option	Description
SMTP (Email) Server Address Settings	Allows you to configure Server IP Address, SMTP Port Number and enable the authentication. In case you are enabling authentication, you need to provide username and password.
Email Addresses	Allows you to configure multiple Email IDs to receive email notifications about system status change. You can send one test email to the configured account from the system.
Alert Categories	Allows you to select multiple alert categories to receive email notifications.

NOTE: Any member iDRAC with system lockdown enabled, that is part of the same group returns an error that the user password was not updated.

Export

Use this section to export the Group Summary to the local system. The information can be exported to a csv file format. It contains data related to each individual system in the group. Export includes the following information in csv format. Server details:

- Health
- Host Name
- iDRAC IPV4 Address
- iDRAC IPV6 Address
- Asset Tag
- Model
- iDRAC Firmware Version
- Last Status Update
- Express Service Code
- iDRAC Connectivity
- Power State
- Operating System
- Service Tag
- Node ID
- iDRAC DNS Name
- BIOS Version
- CPU Details
- System Memory(MB)
- Location Details

NOTE: In case, you are using Internet Explorer, disable the Enhanced Security settings to successfully download the csv file.

Discovered Servers View

After creating the local group, iDRAC group manager notifies all other iDRACs on the local network that a new group has been created. For iDRACs to be displayed under Discovered Servers, group manager feature should be enabled in each iDRAC. Discovered Servers View displays the list of the iDRACs detected on the same network, which can be part of any group. If an iDRAC does not show up in the discovered systems list then the user must logon to the specific iDRAC and join the group. The iDRAC that created the group will be shown as the only member in the essentials view until more iDRACs are joined to the group.

NOTE: Discovered servers view at the Group manager console allows you to onboard one or more servers listed in the view in to that group. The progress of the activity can be tracked from **GroupManager** > **Jobs**. Alternatively you can logon to



the iDRAC and select the group you would like to onboard from the drop down list to join that group. You can access the GroupManager welcome screen from iDRAC index page.

Table 38. Group onboard options

Option	Description
Onboard and Change Login	Select a specific row and select the Onboard and Change Login option to get the newly discovered systems to the group. You must provide the admin logon credentials for the new systems to join the group. If the system has the default password, you need to change it while onboarding it to a group. Group onboarding allows you to apply the same group alert settings to the new systems.
Ignore	Allows you to ignore the systems from the discovered servers list, in case you do not want to add them in any group.
Un-Ignore	Allows you to select the systems that you would like to reinstate in the discovered servers list.
Rescan	Allows you to scan and generate the list of discovered servers at any time.

Jobs View

Jobs view allows the user to track the progress of a group job, helps with simple recovery steps to correct connectivity induced failures. It also shows the history of the last group actions that were performed as an audit log. The user can use the jobs view to track the progress of the action across the group or to cancel an action that is schedule to occur in the future. The Jobs view allows the user to view the status of the last 50 jobs that have been run and any success or failures that has occurred.

Table 39. Jobs View

Option	Description
Status	Shows the job's status and the state of the ongoing job.
dop	Displays the Job's name.
ID	Displays the Job's ID.
Start Time	Displays the start time.
End Time	Displays the end time.
Actions	 Cancel — A scheduled job can be cancelled, before it moves to running state. A running job can be stopped by using the stop button. Rerun — Allows the user to rerun the job in case the job is in failure state. Remove — Allows the user to remove the completed old jobs.
Export	You can export the group job information to the local system for future references. The jobs list can be exported to csv file format. It contains data related to individual job.

NOTE: For each job entry, the list of systems provide details up to 100 systems. Each system entry contains Hostname, Service Tag, Member Job Status and Message if in case the job failed.

All Group actions that create jobs are performed on all the group members with immediate effect. You can perform the following tasks:

- Add/Edit/Remove users
- Configure email alerts



• Change group passcode and name

NOTE: Group jobs complete quickly as long as all members are online and accessible. It may take 10 minutes from job start to job complete. A job will wait and retry for up to 10 hours for the systems that are not accessible.

() NOTE: While an onboarding job is running no other Job can be scheduled. Jobs include:

- Add New User
- Change User Password
- Delete User
- Configure Alerts
- Onboard additional systems
- Change Group Passcode
- Change Group Name

Attempting to invoke another Job while an Onboarding task is active, consequences GMGR0039 error code. Once the onboarding task has made its first attempt to onboard all the new systems, Jobs can be created at any point in time.

Jobs Export

You can export the log to the local system for further references. The jobs list can be exported to a csv file format. It contains all the data related to each job.

(i) NOTE: Exported CSV files are available only in English.

Group Information Panel

Group Information panel in the top right of group manager summary view shows a consolidated group summary. Current group configuration can be edited from the Group Settings page accessible by clicking Group Settings button. It shows how many systems are there in the group. It also provides the information about the Primary and the Secondary controller of the Group.

Group Settings

Group settings page provides a listing of selected group attributes.

Table 40. Group setting attributes

Group Attribute	Description
Group Name	Displays the name of that Group.
Number of Systems	Displays the total number of systems in that group.
Created on	Displays the time stamp details.
Created by	Displays the details of Group admin.
Controlling System	Displays the Service Tag of the system, that acts as the controlling system and coordinates the group management tasks.
Backup System	Displays the Service Tag of the system, that acts as the backup system. In case the controlling system is unavailable, it takes up roles of the controlling system.

Allows the user to perform actions listed on the table below on the group. A group configuration job would be created for these actions (change group name, change group passcode, remove the members and delete the group). The status of group job can be viewed or modified from **GroupManager** > **Jobs** page.



Table 41. Group setting actions

Actions	Description
Change Name	Allows you to change the Current Group Name with a New Group Name .
Change Passcode	Allows you to change the existing group password by entering a New Group Passcode and validating that password by Reenter New Group Passcode .
Remove Systems	Allows you to remove multiple systems from the group at a time.
Delete Group	Allows you to delete the group. To use any feature of group manager, the user should have administrator privileges. Any pending jobs will be stopped in case the group is deleted.

Actions on a selected Server

On the Summary page, you can double click on a row to launch iDRAC for that server through a single sign on redirect. Ensure to turn off the popup blocker in the browser settings. You can perform following actions on the selected server by clicking appropriate item from the **More Actions** drop down list.

Table 42. Actions on a selected Server

Option	Description
Graceful Shutdown	Shuts down the operating system and powers off the system.
Cold Reboot	Powers off, then reboots the system.
Virtual Console	Launches Virtual Console with single sign on a new browser window. (i) NOTE: Disable Popup blocker from the browser to use this functionality.

Group Manager Single Sign On

All iDRACs in the group trust each other based on the shared passcode secret and shared group name. As a result an administrator user at a group member IDRAC is grant administrator level privileges at any group member iDRAC when accessed through Group Manager web interface single sign on. iDRACs logs <user>-<SVCTAG> as the user that logged on into peer members. <SVCTAG> is the service tag of the iDRAC where the user first logged in.

Group Manager Concepts — Controlling System

- Automatically selected by default the first iDRAC configured for Group Manager.
- Provides Group Manager GUI workflow.
- Keeps track of all members.
- Coordinates tasks.
- If a user logs in to any member and clicks on Open Group Manager the browser will be redirected to the primary controller.

Group Manager Concepts — Backup System

- Primary controller automatically selects a secondary controller to take over if the primary goes offline for an extended period of time (10 mins or more).
- If both primary and secondary goes offline for an extended duration (for more than 14 mins) a new primary and secondary controller gets elected.
- Keeps a copy of the group manager cache of all the groups members and tasks.
- The controlling system and backup system are automatically determined by group manager.



• No user configuration or involvement required.

iDRAC Group Firmware Update

For iDRAC group firmware update, from the DUP file from a local directory, perform the following steps:

- 1. Access group manager console essential view and click Update iDRAC Firmware under summary view.
- 2. From the firmware update dialog box displayed, browse and select the local iDRAC DUP file to be installed. Click Upload.
- **3.** File is uploaded to iDRAC and verified for integrity.
- 4. Confirm the firmware update. Group iDRAC firmware update job is scheduled for immediate execution. If Group Manager has other group jobs running, then it is executed after the previous job is completed.
- 5. You can track iDRAC update job execution from group jobs view.

(i) NOTE: This feature is only supported on iDRAC version 3.50.50.50 and above.



Managing logs

iDRAC provides Lifecycle log that contains events related to system, storage devices, network devices, firmware updates, configuration changes, license messages, and so on. However, the system events are also available as a separate log called System Event Log (SEL). The lifecycle log is accessible through iDRAC Web interface, RACADM, and WSMan interface.

When the size of the lifecycle log reaches 800 KB, the logs are compressed and archived. You can only view the non-archived log entries, and apply filters and comments to non-archived logs. To view the archived logs, you must export the entire lifecycle log to a location on your system.

Topics:

- Viewing System Event Log
- Viewing Lifecycle log
- Exporting Lifecycle Controller logs
- Adding work notes
- Configuring remote system logging

Viewing System Event Log

When a system event occurs on a managed system, it is recorded in the System Event Log (SEL). The same SEL entry is also available in the LC log.

(i) NOTE: SEL and LC logs may have mismatch in timestamp when iDRAC is rebooting.

Viewing System Event Log using web interface

To view the SEL, in iDRAC Web interface, go to Maintenance > System Event Log.

The **System Event Log** page displays a system health indicator, a time stamp, and a description for each event logged. For more information, see the *iDRAC Online Help*.

Click Save As to save the SEL to a location of your choice.

NOTE: If you are using Internet Explorer and if there is a problem when saving, download the Cumulative Security Update for Internet Explorer. You can download it from the Microsoft Support website at **support.microsoft.com**.

To clear the logs, click Clear Log.

(i) NOTE: Clear Log only appears if you have Clear Logs permission.

After the SEL is cleared, an entry is logged in the Lifecycle Controller log. The log entry includes the user name and the IP address from where the SEL was cleared.

Viewing System Event Log using RACADM

To view the SEL: racadm getsel <options> If no arguments are specified, the entire log is displayed. To display the number of SEL entries: racadm getsel -i To clear the SEL entries: racadm clrsel For more information, see *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Managing logs 193



Viewing System Event Log using iDRAC settings utility

You can view the total number of records in the System Event Log (SEL) using the iDRAC Settings Utility and clear the logs. To do this:

- In the iDRAC Settings Utility, go to System Event Log. The iDRAC Settings.System Event Log displays the Total Number of Records.
- 2. To clear the records, select Yes. Else, select No.
- 3. To view the system events, click **Display System Event Log**.
- 4. Click Back, click Finish, and then click Yes.

Viewing Lifecycle log

Lifecycle Controller logs provide the history of changes related to components installed on a managed system. You can also add work notes to each log entry.

The following events and activities are logged:

- All
- System Health System Health category represents all the alerts that are related to hardware within the system chassis.
- Storage Storage Health category represents alerts that are related to the storage subsystem.
- Updates Update category represents alerts that are generated due to firmware/Driver upgrades/downgrades.
- Audit Audit category represents the audit log.
- Configuration Configuration category represents alerts that are related to hardware, firmware and software configuration changes.
- Work Notes

When you log in to or log out of iDRAC using any of the following interfaces, the log-in, log-out, or login failure events are recorded in the Lifecycle logs:

- SSH
- Web interface
- RACADM
- Redfish
- IPMI over LAN
- Serial
- Virtual console
- Virtual media

You can view and filter logs based on the category and severity level. You can also export and add a work note to a log event.

(i) NOTE: Lifecycle logs for Personality Mode change is generated only during the warm boot of the host.

If you initiate configuration jobs using RACADM CLI or iDRAC web interface, the Lifecycle log contains information about the user, interface used, and the IP address of the system from which you initiate the job.

NOTE: On MX platform, Lifecycle Controller logs multiple job IDs for configuration or installation jobs created using OME - Modular. For more information on the jobs performed, see the OME - Modular logs.

Viewing Lifecycle log using web interface

To view the Lifecycle Logs, click **Maintenance** > **Lifecycle Log**. The **Lifecycle Log** page is displayed. For more information about the options, see the *iDRAC Online Help*.

Filtering Lifecycle logs

You can filter logs based on category, severity, keyword, or date range.

To filter the lifecycle logs:

1. In the Lifecycle Log page, under the Log Filter section, do any or all of the following:



- Select the Log Type from the drop-down list.
- Select the severity level from the **Severity** drop-down list.
- Enter a keyword.
- Specify the date range.
- 2. Click Apply.

The filtered log entries are displayed in Log Results.

Adding comments to Lifecycle logs

To add comments to the Lifecycle logs:

- In the Lifecycle Log page, click the + icon for the required log entry. The Message ID details are displayed.
- Enter the comments for the log entry in the Comment box. The comments are displayed in the Comment box.

Viewing Lifecycle log using RACADM

To view Lifecycle logs, use the lclog command.

For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Exporting Lifecycle Controller logs

You can export the entire Lifecycle Controller log (active and archived entries) in a single zipped XML file to a network share or to the local system. The zipped XML file extension is .xml.gz. The file entries are ordered sequentially based on their sequence numbers, ordered from the lowest sequence number to the highest.

Exporting Lifecycle Controller logs using web interface

To export the Lifecycle Controller logs using the Web interface:

- 1. In the Lifecycle Log page, click Export.
- 2. Select any of the following options:
 - Network Export the Lifecycle Controller logs to a shared location on the network.
 - Local Export the Lifecycle Controller logs to a location on the local system.
 - **NOTE:** While specifying the network share settings, it is recommended to avoid special characters for user name and password or percent encode the special characters.

For information about the fields, see the iDRAC Online Help.

3. Click Export to export the log to the specified location.

Exporting Lifecycle Controller logs using RACADM

To export the Lifecycle Controller logs, use the lclog export command.

For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Adding work notes

Each user who logs in to iDRAC can add work notes and this is stored in the lifecycle log as an event. You must have iDRAC logs privilege to add work notes. A maximum of 255 characters are supported for each new work note.

(i) NOTE: You cannot delete a work note.

To add a work note:



- In the iDRAC Web interface, go to Dashboard > Notes > add note. The Work Notes page is displayed.
- 2. Under Work Notes, enter the text in the blank text box.

(i) NOTE: It is recommended not to use too many special characters.

3. Click Save.

The work note is added to the log. For more information, see the *iDRAC Online Help*.

Configuring remote system logging

You can send lifecycle logs to a remote system. Before doing this, make sure that:

- There is network connectivity between iDRAC and the remote system.
- The remote system and iDRAC is on the same network.

Configuring remote system logging using web interface

To configure the remote syslog server settings:

- In the iDRAC Web interface, go to Configuration > System Settings > Remote Syslog Settings. The Remote Syslog Settings page is displayed
- 2. Enable remote syslog, specify the server address, and the port number. For information about the options, see the *iDRAC Online Help*.
- 3. Click Apply.

The settings are saved. All logs written to the lifecycle log are also simultaneously written to configured remote server(s).

Configuring remote system logging using RACADM

To configure the remote system-logging settings, use the set command with the objects in the iDRAC.SysLog group.

For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.



Monitoring and managing power in iDRAC

You can use iDRAC to monitor and manage the power requirements of the managed system. This helps to protect the system from power outages by appropriately distributing and regulating the power consumption on the system.

The key features are:

- **Power Monitoring** View the power status, history of power measurements, the current averages, peaks, and so on for the managed system.
- **Power Capping** View and set the power cap for the managed system, including displaying the minimum and maximum potential power consumption. This is a licensed feature.
- **Power Control** Enables you to remotely perform power control operations (such as, power on, power off, system reset, power cycle, and graceful shutdown) on the managed system.
- **Power Supply Options** Configure the power supply options such as redundancy policy, hot spare, and power factor correction.

Topics:

- Monitoring power
- Setting warning threshold for power consumption
- Executing power control operations
- Power capping
- Configuring power supply options
- Enabling or disabling power button
- Multi-Vector Cooling

Monitoring power

iDRAC monitors the power consumption in the system continuously and displays the following power values:

- Power consumption warning and critical thresholds.
- Cumulative power, peak power, and peak amperage values.
- Power consumption over the last hour, last day or last week.
- Average, minimum, and maximum power consumption.
- Historical peak values and peak timestamps.
- Peak headroom and instantaneous headroom values (for rack and tower servers).

NOTE: The histogram for the system power consumption trend (hourly, daily, weekly) is maintained only while iDRAC is running. If iDRAC is restarted, the existing power consumption data is lost and the histogram is restarted.

(i) NOTE: After iDRAC firmware update or reset, the power consumption graph will be wiped / reset.

Monitoring performance index of CPU, memory, and input output modules using web interface

To monitor the performance index of CPU, memory, and I/O modules, in the iDRAC web interface, go to **System** > **Performance**.

- System Performance section Displays the current reading and the warning reading for CPU, Memory and I/O utilization index, and system level CUPS index in a graphical view.
- System Performance Historical Data section:
 - Provides the statistics for CPU, memory, IO utilization, and the system level CUPS index. If the host system is powered off, then the graph displays the power off line below 0 percent.

Monitoring and managing power in iDRAC 197



- You can reset the peak utilization for a particular sensor. Click **Reset Historical Peak**. You must have Configure privilege to reset the peak value.
- Performance Metrics section:
 - Displays status and present reading
 - Displays or specifies the warning threshold utilization limit. You must have server configure privilege to set the threshold values.

For information about the displayed properties, see the *iDRAC Online Help*.

Monitoring performance index for of CPU, memory, and input output modules using RACADM

Use the **SystemPerfStatistics** sub command to monitor performance index for CPU, memory, and I/O modules. For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Setting warning threshold for power consumption

You can set the warning threshold value for the power consumption sensor in the rack and tower systems. The warning/critical power threshold for rack and tower systems may change, after the system is power-cycled, based on PSU capacity and redundancy policy. However, the warning threshold must not exceed the critical threshold even if Power Supply Unit capacity of the redundancy policy is changed.

The warning power threshold for blade systems is set to power allocation of CMC(for non-MX platforms) or OME Modular (for MX platforms).

If reset to default action is performed, the power thresholds will be set to default.

You must have Configure user privilege to set the warning threshold value for power consumption sensor.

(i) NOTE: The Warning Threshold value is reset to the default value after performing a racreset or an iDRAC update.

Setting warning threshold for power consumption using web interface

- 1. In the iDRAC Web interface, go to System > Overview > Present Power Reading and Thresholds.
- 2. In the **Present Power Reading and Thresholds** section, click **Edit Warning Threshold**. The **Edit Warning Threshold** page is displayed.
- 3. In the Warning Threshold column, enter the value in Watts or BTU/hr. The values must be lower than the Failure Threshold values. The values are rounded off to the nearest value that is divisible by 14. If you enter Watts, the system automatically calculates and displays the BTU/hr value. Similarly, if you enter BTU/hr, the value for Watts is displayed.
- 4. Click Save. The values are configured.

Executing power control operations

iDRAC enables you to remotely perform a power-on, power off, reset, graceful shutdown, Non-Masking Interrupt (NMI), or power cycle using the Web interface or RACADM.

You can also perform these operations using Lifecycle Controller Remote Services or WSMan. For more information, see the Lifecycle Controller Remote Services Quick Start Guide available at https://www.dell.com/idracmanuals and the Dell Power State Management Profile document available at https://www.dell.com/support.

Server power-control operations initiated from iDRAC are independent of the power-button behavior configured in the BIOS. You can use the PushPowerButton function to gracefully shut down the system, or power it on, even if the BIOS is configured to do nothing when the physical power button is pressed.

198 Monitoring and managing power in iDRAC



Executing power control operations using web interface

To perform power control operations:

- 1. In iDRAC web interface, go to **Configuration > Power Management > Power Control**. The **Power Control** options are displayed.
- **2.** Select the required power operation:
 - Power On System
 - Power Off System
 - NMI (Non-Masking Interrupt)
 - Graceful Shutdown
 - Reset System (warm boot)
 - Power Cycle System (cold boot)
- 3. Click **Apply**. For more information, see the *iDRAC Online Help*.

Executing power control operations using RACADM

To perform power actions, use the **serveraction** command.

For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Power capping

You can view the power threshold limits that covers the range of AC and DC power consumption that a system under heavy workload presents to the datacenter. This is a licensed feature.

Power capping in Blade servers

Before a blade server turns on, based on limited hardware inventory, iDRAC provides the power requirements of the blade server to the chassis manager. If the power consumption increases over time and if the server consumes power to its maximum allocation, iDRAC requests CMC (for non-MX platforms) or OME Modular (for MX platforms) to increase the maximum potential power. This results in an increase in the power delivery, however, the power delivery does not reduce if the consumption decreases.

After the system is powered on and initialized, iDRAC calculates a new power requirement based on the actual hardware configuration. The system stays powered on even if the CMC (not for MX platforms) or OME Modular (not for MX platforms) fails to allocate new power request.

CMC or OME Modular reclaims any unused power from lower priority servers and allocates that power to a higher-priority infrastructure module or a server.

Viewing and configuring power cap policy

When power cap policy is enabled, it enforces a user-defined power limits on the system. If power-capping is not enabled, the default hardware power-protection policy is used. This power-protection policy is independent of the user-defined policy. The system performance is dynamically adjusted to maintain power consumption close to the specified threshold.

Actual power consumption depends on the workload. It may momentarily exceed the threshold until performance adjustments are completed. For example, consider a system that has a minimum and maximum Potential Power Consumption values of 500 W and 700 W respectively. You can specify a Power Budget Threshold to reduce consumption to 525 W. When this power budget is configured, the performance of the system is dynamically adjusted to maintain power consumption of 525 W or less.

If you set a very low power cap or if the ambient temperature is unusually high, power consumption may temporarily exceed the power-cap while the system is powering up or being reset.

If the power cap value is set lower than the minimum recommended threshold, iDRAC may not be able maintain the requested power cap.

You can specify the value in Watts, BTU/hr, or as a percentage of the recommended maximum power limit.



When setting the power cap threshold in BTU/hr, the conversion to Watts is rounded off to the nearest integer. When the power cap threshold are read from the system, the Watts to BTU/hr conversion is also rounded off. Because of the rounding off, the actual values may slightly differ.

Configuring power cap policy using web interface

To view and configure the power policies:

- 1. In iDRAC Web interface, go to **Configuration** > **Power Management** > **Power Cap Policy**. The current power policy limit is displayed under the **Power Cap Limits** section.
- 2. Select Enable under Power Cap.
- 3. Under **Power Cap Limits** section, enter the power limit within recommended range in Watts and BTU/hr or the maximum % of recommended system limit.
- 4. Click Apply to apply the values.

Configuring power cap policy using RACADM

To view and configure the current power cap values, use the following objects with the set command:

- System.Power.Cap.Enable
- System.Power.Cap.Watts
- System.Power.Cap.Btuhr
- System.Power.Cap.Percent

For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Configuring power cap policy using iDRAC settings utility

To view and configure power policies:

1. In iDRAC Settings utility, go to Power Configuration.

(i) NOTE: The Power Configuration link is available only if the server power supply unit supports power monitoring.

The iDRAC Settings Power Configuration page is displayed.

- 2. Select Enabled to enable the Power Cap Policy Else, select Disabled.
- **3.** Use the recommended settings, or under **User Defined Power Cap Policy**, enter the required limits. For more information about the options, see the *iDRAC Settings Utility Online Help*.
- Click Back, click Finish, and then click Yes. The power cap values are configured.

Configuring power supply options

You can configure the power supply options such as redundancy policy, hot spare, and power factor correction.

Hot spare is a power supply feature that configures redundant Power Supply Units (PSUs) to turn off depending on the server load. This allows the remaining PSUs to operate at a higher load and efficiency. This requires PSUs that support this feature, so that it quickly powers ON when needed.

In a two PSU system, either PSU1 or PSU2 can be configured as the primary PSU.

After Hot Spare is enabled, PSUs can become active or go to sleep based on load. If Hot Spare is enabled, asymmetric electrical current sharing between the two PSUs is enabled. One PSU is *awake* and provides the majority of the current; the other PSU is in sleep mode and provides a small amount of the current. This is often called 1 + 0 with two PSUs and hot spare enabled. If all PSU-1s are on Circuit-A and all PSU-2s are on Circuit-B, then with hot spare enabled (default hot spare factory configuration), Circuit-B has much less load and triggers the warnings. If hot spare is disabled, the electrical current sharing is 50-50 between the two PSUs, the Circuit-A and Circuit-B normally has the same load.

Power factor is the ratio of real power consumed to the apparent power. When power factor correction is enabled, the server consumes a small amount of power when the host is OFF. By default, power factor correction is enabled when the server is shipped from the factory.

200 Monitoring and managing power in iDRAC


Configuring power supply options using web interface

To configure the power supply options:

- 1. In iDRAC Web interface, go to Configuration > Power Management > Power Configuration.
- 2. Under **Power Redundancy Policy**, select the required options. For more information, see *iDRAC Online Help*.
- 3. Click **Apply**. The power supply options are configured.

Configuring power supply options using RACADM

To configure the power supply options, use the following objects with the get/set command:

- System.Power.RedundancyPolicy
- System.Power.Hotspare.Enable
- System.Power.Hotspare.PrimaryPSU
- System.Power.PFC.Enable

For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Configuring power supply options using iDRAC settings utility

To configure the power supply options:

1. In iDRAC Settings utility, go to **Power Configuration**.

(i) NOTE: The Power Configuration link is available only if the server power supply unit supports power monitoring.

The iDRAC Settings Power Configuration page is displayed.

- 2. Under Power Supply Options:
 - Enable or disable power supply redundancy.
 - Enable or disable hot spare.
 - Set the primary power supply unit.
 - Enable or disable power factor correction. For more information about the options, see the *iDRAC Settings Utility Online Help*.
- Click Back, click Finish, and then click Yes. The power supply options are configured.

Enabling or disabling power button

To enable or disable the power button on the managed system:

- 1. In iDRAC Settings utility, go to Front Panel Security.
- The **iDRAC Settings Front Panel Security** page is displayed.
- 2. Select Enabled to enable the power button or Disabled to disable it.
- **3.** Click **Back**, click **Finish**, and then click **Yes**. The settings are saved.

Multi-Vector Cooling

Multi-Vector Cooling implements multi-prong approach to Thermal Controls in Dell EMC Server Platforms. You can configure multi-vector cooling options through iDRAC web interface by navigating to **Configuration** > **System Settings** > **Hardware Settings** > **Fan Configuration**. It includes (but not limited to):

- Large set of sensors (thermal, power, inventory etc.) that allows accurate interpretation of real-time system thermal state at various locations within the server. It displays only a small subset of sensors that are relevant to users need based on the configuration.
- Intelligent and adaptive closed loop control algorithm optimizes fan response to maintain component temperatures. It also conserves fan power, airflow consumption, and acoustics.

Monitoring and managing power in iDRAC 201



- Using fan zone mapping, cooling can be initiated for the components when it requires. Thus, it results maximum performance without compromising the efficiency of power utilization.
- Accurate representation of slot by slot PCIe airflow in terms of LFM metric (Linear Feet per Minute an accepted industry standard on how PCIe card airflow requirement is specified). Display of this metric in various iDRAC interfaces allows user to:
 - 1. know the maximum LFM capability of each slot within the server.
 - 2. know what approach is being taken for PCIe cooling for each slot (airflow controlled, temperature controlled).
 - 3. know the minimum LFM being delivered to a slot, if the card is a 3rd Party Card (user defined custom card).
 - **4.** dial in custom minimum LFM value for the 3rd Party Card allowing more accurate definition of the card cooling needs for which the user is better aware of through their custom card specification.
- Displays real-time system airflow metric (CFM, cubic feet per minute) in various iDRAC interfaces to the user to enable datacenter airflow balancing based on aggregation of per server CFM consumption.
- Allows custom thermal settings like Thermal Profiles (Maximum Performance vs. Maximum Performance per Watt, Sound
 - Cap), custom fan speed options (minimum fan speed, fan speed offsets) and custom Exhaust Temperature settings.1. Most of these settings allow additional cooling over the baseline cooling generated by thermal algorithms and do not allow fan speeds to go below system cooling requirements.
 - () **NOTE:** One exception to above statement is for fan speeds that are added for 3rd Party PCIe cards. The thermal algorithm provision airflow for 3rd party cards may be more or less than the actual card cooling needs and customer may fine tune the response for the card by entering the LFM corresponding to the 3rd Party Card.
- 2. Custom Exhaust Temperature option limits exhaust temperature to customer desired settings.
 - **NOTE:** It is important to note that with certain configurations and workloads, it may not be physically possible to reduce exhaust below a desired set point (e.g. Custom exhaust setting of 45C with a high inlet temp {e.g. 30C} and a loaded config {high system power consumption, low airflow}).
- **3.** Sound Cap option is new in the 14th generation of PowerEdge server. It limits CPU power consumption and controls fan speed and acoustical ceiling. This is unique for acoustical deployments and may result in reduced system performance.
- System layout and design enables increased airflow capability (by allowing high power) and dense system configurations. It provides less system restrictions and increased feature density.
 - 1. Streamlined airflow permits efficient airflow to fan power consumption ratio.
- Custom fans are designed for higher efficiency, better performance, longer life and less vibration. It also delivers better acoustics outcome.
 - **1.** Fans are capable of long life (in general it may run for more than 5 years), even if it runs at full speed all the time.
- Custom heat-sinks are designed for optimize component cooling at minimum (required) airflow yet supports high performance CPUs.



iDRAC Direct Updates

iDRAC provides out of band ability to update the firmware of various components of a PowerEdge server. iDRAC direct update helps in eliminating staged jobs during updates.

iDRAC used to have staged updates to initiate firmware update of the components. From this release, Direct updates have been applied to PSU and Backplane. With the use of Direct Updates and Backplane can have quicker updates. In case of PSU, one reboot (for initializing the updates) is avoided and the update can happen in single reboot.

With Direct update feature in iDRAC, you can eliminate the first reboot to initiate the updates. The second reboot will be controlled by the device itself and iDRAC notifies the user if there is need for a separate reset via job status.

iDRAC Direct Updates 203



Inventorying, monitoring, and configuring network devices

You can inventory, monitor, and configure the following network devices:

- Network Interface Cards (NICs)
- Converged Network Adapters (CNAs)
- LAN On Motherboards (LOMs)
- Network Daughter Cards (NDCs)
- Mezzanine cards (only for blade servers)

Before you disable NPAR or an individual partition on CNA devices, ensure that you clear all I/O identity attributes (Example: IP address, virtual addresses, initiator, and storage targets) and partition-level attributes (Example: Bandwidth allocation). You can disable a partition either by changing the VirtualizationMode attribute setting to NPAR or by disabling all personalities on a partition.

Depending on the type of installed CNA device, the settings of partition attributes may not be retained from the last time the partition was active. Set all I/O identity attributes and partition-related attributes when enabling a partition. You can enable a partition by either changing the VirtualizationMode attribute setting to NPAR or by enabling a personality (Example: NicMode) on the partition.

Topics:

- Inventorying and monitoring network devices
- Inventorying and monitoring FC HBA devices
- Inventorying and monitoring SFP Transceiver devices
- Telemetry Streaming
- Serial Data Capture
- Dynamic configuration of virtual addresses, initiator, and storage target settings

Inventorying and monitoring network devices

You can remotely monitor the health and view the inventory of the network devices in the managed system.

For each device, you can view the following information of the ports and enabled partitions:

- Link Status
- Properties
- Settings and Capabilities
- Receive and Transmit Statistics
- iSCSI, FCoE initiator, and target information

Monitoring network devices using web interface

To view the network device information using Web interface, go to **System** > **Overview** > **Network Devices**. The **Network Devices** page is displayed. For more information about the displayed properties, see *iDRAC Online Help*.

Monitoring network devices using RACADM

To view information about network devices, use the hwinventory and nicstatistics commands.

For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Additional properties may be displayed when using RACADM or WSMan in addition to the properties displayed in the iDRAC web interface.



Connection View

Manually checking and troubleshooting the servers' networking connections is unmanageable in a datacenter environment. iDRAC9 streamlines the job with iDRAC Connection View. This feature allows you to remotely check and troubleshoot network connections from the same centralized GUI that you are using for deploying, updating, monitoring, and maintaining the servers. Connection View in iDRAC9 provide details of the physical mapping of switch ports to server's network ports and iDRAC (integrated Dell Remote Access Controller) dedicated port connections. All supported network cards are visible in Connection View, irrespective of the brand.

Instead of manually checking and troubleshooting the server's networking connections, you can view and manage network cable connections remotely.

The Connection View provides the information of the switch ports which are connected to the server ports, and iDRAC dedicated port. The server network ports include those on PowerEdge LOM, NDC, Mezz cards, and PCIe add-in cards.

To View network devices connection view, navigate to **System** > **Overview** > **Network Device** > **Connection View** to view the Connection View.

Also, you can click **iDRAC Settings** > **Connectivity** > **Network** > **Common Settings** > **Connection View** to enable or disable the connection view.

Connection View can be explored with racadm SwitchConnection View command and it can also be viewed with command.

Field or Option	Description
Enabled	Select Enabled , to enable Connection View. By default the Enabled option is selected.
State	Displays Enabled , if you enable the connection view option from the Connection View from iDRAC settings.
Switch Connection ID	Displays the LLDP chassis ID of the switch through which the device port is connected.
Switch Port Connection ID	Displays the LLDP port ID of the switch port to which the device port is connected.

() NOTE: Switch Connection ID and Switch Port Connection ID are available once the Connection View is enabled and the Link is connected. The associated network card needs to be compatible with the Connection View. Only users with iDRAC Configure privilege can modify the Connection View settings.

From iDRAC9 4.00.00.00 and later versions, iDRAC supports to send standard LLDP packets to external switches. This provides options to discover iDRACs on the network. iDRAC sends two types of LLDP packets to the outbound network:

• **Topology LLDP** - In this feature, the LLDP packet goes through all the supported server NIC ports so that an external switch can locate the originating server, NDC port[NIC FQDD], IOM location in the chassis, blade chassis service tag etc. From iDRAC9 4.00.00.00 and later versions, Topology LLDP is available as an option for all PowerEdge servers. The LLDP packets contain server network device connectivity information and are used by I/O modules and external switches to update their configuration.

() NOTE:

- The Topology LLDP must be enabled for the MX chassis configuration to function properly.
- The Topology LLDP is not supported on 1GbE controllers and select 10GbE controllers (Intel X520, QLogic 578xx).
- **Discovery LLDP** In this feature, the LLDP packet goes only through the active iDRAC NIC port in use (dedicated NIC or shared LOM), so an adjacent switch can locate iDRAC connection port in the switch. Discovery LLDP is specific only to the active iDRAC network port and not will not be seen in all the Network ports in the server. Discovery LLDP will have some details of the idrac like IP address, MAC address, service tag etc., so that a switch can automatically discover iDRAC devices connected to it and some data of iDRAC.

(i) NOTE: If Virtual MAC Address is cleared on a port/partition, then Virtual MAC Address would be same as MAC Address.

To enable or disable the Topology LLDP, navigate to **iDRAC Settings** > **Connectivity** > **Network** > **Common Settings** > **Topology LLDP** to enable or disable the Topology LLDP. By default, it is enabled for MX servers and disabled for all other servers.

To enable or disable the iDrac Discovery LLDP, navigate to **iDRAC Settings** > **Connectivity** > **Network** > **Common Settings** > **iDrac Discovery LLDP**. By default, the Enabled option is selected.

LLDP packet originated from idrac can be viewed from switch using the command: show lldp neighbors.



Refresh Connection View

Use Refresh Connection View to get the latest information of Switch Connection ID and Switch Port Connection ID.

() NOTE: If iDRAC has switch connection and switch port connection information for server network port or iDRAC network port and due to some reason, the switch connection and switch port connection information is not refreshed for 5min, then the switch connection and switch port connection information is shown as stale (last known good data) data for all user interfaces. In the UI, you see yellow bang which is a natural representation and it does not indicate any warning.

Connection View Possible Values

Possible Connection View Data	Description
Feature Disabled	Connection view feature is disabled, to view the connection view data enable the feature.
No Link	Indicates that the link associated with network controller port is down.
Not Available	LLDP is not enabled on the switch. Check whether LLDP is enabled on the switch port.
Not Supported	Network controller does not support Connection view feature.
Stale Data	Last known good data, either the Network controller port link is down or the system is powered off. Use the refresh option to refresh the connection view details to get the latest data.
Valid Data	Displays the Valid Switch Connection ID and the Switch Port Connection ID information.

Connection View Supported Network Controllers

Following cards or controllers support Connection View feature.

Manufacturer	Туре
Broadcom	 57414 rNDC 25GE 57416/5720 rNDC 10GbE 57412/5720 rNDC 10GbE 57414 PCIe FH/LP 25GE 57412 PCIe FH/LP 10GbE 57416 PCIe FH/LP 10GbE
Intel	 X710 bNDC 10Gb X710 DP PCIe 10Gb X710 QP PCIe 10Gb X710 + I350 rNDC 10Gb+1Gb X710 rNDC 10Gb X710 bNDC 10Gb XL710 bNDC 10Gb XL710 PCIe 40Gb XL710 OCP Mezz 10Gb X710 PCIe 10Gb
Mellanox	 MT27710 rNDC 40Gb MT27710 PCIe 40Gb MT27700 PCIe 100Gb
QLogic	 QL41162 PCIe 10GE 2P QL41112 PCIe 10GE 2P QL41262 PCIe 25GE 2P



Inventorying and monitoring FC HBA devices

You can remotely monitor the health and view the inventory of the Fibre Channel Host Bus Adapters (FC HBA) devices in the managed system. The Emulex and QLogic FC HBAs are supported. For each FC HBA device, you can view the following information for the ports:

- Link Status and Information
- Port Properties
- Receive and Transmit Statistics

(i) NOTE: Emulex FC8 HBAs are not supported.

Monitoring FC HBA devices using web interface

To view the FC HBA device information using Web interface, go to **System** > **Overview** > **Network Devices** > **Fibre Channel**. For more information about the displayed properties, see *iDRAC Online Help*.

The page name also displays the slot number where the FC HBA device is available and the type of FC HBA device.

Monitoring FC HBA devices using RACADM

To view the FC HBA device information using RACADM, use the hwinventory command.

For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Inventorying and monitoring SFP Transceiver devices

You can remotely monitor the health and view the inventory of SFP transceiver devices connected to the system. Following are the supported transceivers:

- SFP
- SFP+
- SFP28
- SFP-DD
- QSFP
- QSFP+
- QSFP28
- QSFP-DD
- Base-T modules
- AOC & DAC cables
- RJ-45 Base-T connected with Ethernet
- Fiber channel
- IB adapter ports

Most useful transceiver information are Serial number and Part number from transceiver EPROM. These would allow to verify the remotely installed transceivers, when troubleshooting connectivity issues. For each SFP Transceiver device, you can view the following information for the ports:

- Vendor Name
- Part Number
- Revision
- Serial Number
- Device Identifier / Type Info
- Cable length (in meter)



Monitoring SFP Transceiver devices using web interface

To view theSFP Transceiver device information using Web interface, go to **System** > **Overview** > **Network Devices** and click on particular device. For more information about the displayed properties, see *iDRAC Online Help*.

The page name also displays the slot number where the transceiver device is available under Port statistics.

Monitoring SFP Transceiver devices using RACADM

To view the SFP Transceiver device information using RACADM, use the hwinventory command.

For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Telemetry Streaming

Telemetry enables users to collect and stream real-time device metrics, events, and data logs from a PowerEdge server to a subscribed external client or server application. Using Telemetry, you can set the type and frequency of reports that needs to be generated.

(i) NOTE: The feature is supported on all the platforms and it requires iDRAC Datacenter license.

Telemetry is one-to-many solution for collecting and streaming the live system data from one or more PowerEdge servers (iDRAC) to a centralized 'Remote Server Monitoring, Analysis, and Alerting service'. The feature also supports on-demand data collection of the data.

The telemetry data includes metrics/inventory and logs/events. The data can be streamed (pushed out) or collected (pulled) from iDRAC to or by remote consumers like Redfish client and Remote Syslog Server. The telemetry data is also provided to the iDRAC SupportAssist data collector on demand. The data collection and report is based on predefined Redfish telemetry metrics, trigger, and report definitions. The telemetry streaming settings can be configured through RACADM, Redfish, and Server Configuration Profile (SCP).

To configure Telemetry, enable or select the required device reports or logs that define the behavior and frequency of data streaming. Go to **Configuration** > **System Settings** page to configure Telemetry. Data streaming is automatic until the Telemetry is disabled.

Туре	Metric Group	Inventory	Sensor	Statistics	Configuration	Metrics
I/O Devices	NICs	No	Yes	Yes	No	No
	FC HBAs	No	Yes	Yes	No	No
Server Devices	CPUs	No	Yes	No	No	Yes
	Memory	No	Yes	No	No	Yes
	Fans	No	Yes	No	No	No
	PSUs	No	No	No	No	Yes
	Sensors	No	Yes	No	No	No
Environmental	Thermal	No	Yes	No	No	Yes
	Power	No	No	Yes	No	Yes
	Performance	No	No	Yes	No	No
Accelerators	GPUs	No	No	Yes	No	Yes

Following table describes the metric reports that can be generated using telemetry:

To know about the field descriptions of Telemetry section, see *iDRAC Online Help*.

() NOTE:

• StorageDiskSMARTDATA is only supported on SSD drives with SAS/SATA bus protocol and behind the BOSS controller.



- StorageSensor data is reported only for the drives in Ready / Online / Non-RAID mode and not behind the BOSS controller.
- NVMeSMARTData is only supported for SSD (PCleSSD / NVMe Express) drives with PCle bus protocol (not behind SWRAID).
- GPGPUStatistics data is only available in specific GPGPU models that support ECC memory capability.
- PSUMetrics is not available on modular platforms.
- Fan Power and PCIe Power Metrics may be displayed as 0 for some platforms.
- CUPS report has been renamed to SystemUsage in 4.40.00.00 release and it's supported on both INTEL and AMD platforms.

Telemetry Workflow:

- 1. Install Datacenter license, if not installed already.
- 2. Configure global Telemetry settings including Enabling the telemetry and Rsyslog server network address and port using RACADM, Redfish, SCP, or iDRAC GUI.
- **3.** Configure the following Telemetry report streaming parameters on the required device report or log using either RACADM or Redfish interface:
 - EnableTelemetry
 - ReportInterval
 - ReportTriggers

(i) NOTE: Enable iDRAC Alerts and Redfish events for the specific hardware for which you need telemetry reports.

- 4. Redfish client makes subscription request to the Redfish EventService on iDRAC.
- 5. iDRAC generates and pushes the metric report or log/event data to the subscribed client when the predefined trigger conditions are met.

Feature Constraints:

- 1. For security reasons, iDRAC supports only HTTPS-based communication to the client.
- 2. For stability reasons, iDRAC supports up to eight subscriptions.
- **3.** Deletion of subscriptions is supported through Redfish interface only, even for the manual deletion by the Admin.

Behavior of Telemetry feature:

- iDRAC generates and pushes (HTTP POST) the Metric Report or log/event data to all the subscribed clients to the destination specified in the subscription when the predefined trigger conditions are met. The clients receive new data only upon successful subscription creation.
- The metric data includes the timestamp in ISO format, UTC time (ends in 'Z'), at the time of data collection from source.
- Clients can terminate a subscription by sending an HTTP DELETE message to the URI of the subscription resource through the Redfish interface.
- If the subscription is deleted either by iDRAC or the client, then iDRAC does not send (HTTP POST) reports. If the number of delivery errors exceeds predefined thresholds, then iDRAC may delete a subscription.
- If a user has Admin privilege, they can delete the subscriptions but only through Redfish interface.
- Client is notified about the termination of a subscription by iDRAC by sending 'Subscription terminated' event as the last message.
- Subscriptions are persistent and can remain even after iDRAC restarts. But, they can be deleted either by performing racresetcfg or LCwipe operations.
- User interfaces like RACADM, Redfish, SCP, and iDRAC display the current status of the client subscriptions.

Serial Data Capture

iDRAC allows you to capture console redirection serial for later retrieval with the use of Serial Data Capture feature. This feature requires iDRAC Datacenter license.

The purpose of Serial Data Capture feature is to capture the system serial data and store it so that the customer can later retrieve it for debugging purpose.

You can enable or disable a serial data capture using RACADM, Redfish, iDRAC interfaces. When this attribute is enabled, iDRAC will capture serial traffic received on Host Serial Device2 irrespective of serial Mux mode settings.

To enable / disable Serial Data Capture using iDRAC GUI, go to **Maintainance** > **Diagnostics** > **Serial Data Logs** page, and check the box to enable or disable.



() NOTE:

- This attribute is persistent over iDRAC reboot.
- Firmware reset to default will disable this feature.
- While Serial Data capture is enabled, the buffer keeps getting appended with recent data. If user disables Serial capture and enables it again, iDRAC starts appending from last update.

The System serial data capture starts when user enables the serial data capture flag from any of the interfaces. If serial data capture is enabled after the system has booted, you have to reboot the system, so BIOS can see the new setting (console redirection Enabled requested by iDRAC) to get the serial data. iDRAC will start the data capture continuously and stores to the shared memory with limit of 512 KB. This buffer will be circular.

() NOTE:

- For this feature to be functional, one must have Login privilege and System control privilege.
- This feature requires iDARC Datacenter license.

Dynamic configuration of virtual addresses, initiator, and storage target settings

You can dynamically view and configure the virtual address, initiator and storage target settings, and apply a persistence policy. It allows the application to apply the settings based on power state changes (that is, operating system restart, warm reset, cold reset, or AC cycle) and also based on persistence policy setting for that power state. This provides more flexibility in deployments that need rapid re-configuration of system workloads to another system.

The virtual addresses are:

- Virtual MAC Address
- Virtual iSCSI MAC Address
- Virtual FIP MAC Address
- Virtual WWN
- Virtual WWPN

NOTE: When you clear the persistence policy, all the virtual addresses are reset to the default permanent address set at the factory.

NOTE: Some cards with the virtual FIP, virtual WWN, and virtual WWPN MAC attributes, the virtual WWN and virtual WWPN MAC attributes are automatically configured when you configure virtual FIP.

Using the IO Identity feature, you can:

- View and configure the virtual addresses for network and fibre channel devices (for example, NIC, CNA, FC HBA).
- Configure the initiator (for iSCSI and FCoE) and storage target settings (for iSCSI, FCoE, and FC).
- Specify persistence or clearance of the configured values over a system AC power loss, cold, and warm system resets.

The values configured for virtual addresses, initiator and storage targets may change based on the way the main power is handled during system reset and whether the NIC, CNA, or FC HBA device has auxiliary power. The persistence of IO identity settings can be achieved based on the policy setting made using iDRAC.

Only if the I/O identity feature is enabled, the persistence policies take effect. Each time the system resets or powers on, the values are persisted or cleared based on the policy settings.

(i) NOTE: After the values are cleared, you cannot re-apply the values before running the configuration job.

Supported cards for IO Identity Optimization

The following table provides the cards that support the I/O Identity Optimization feature.

Table 43. Supported cards for I/O Identity Optimization

Manufacturer	Туре
Broadcom	• 5719 Mezz 1GB
	• 5720 PCle 1 GB



Table 43. Supported cards for I/O Identity Optimization (continued)

Manufacturer	Туре
	 5720 bNDC 1 GB 5720 rNDC 1 GB 57414 PCle 25GbE
Intel	 i350 DP FH PCIe 1GB i350 QP PCIe 1GB i350 QP rNDC 1GB i350 Mezz 1GB i350 bNDC 1GB x520 PCIe 10GB x520 bNDC 10GB x520 Mezz 10GB x520 Mezz 10GB x520 + i350 rNDC 10GB+1GB X710 bNDC 10GB X710 PCIe 10 GB X710 PCIe 10 GB X710 rNDC 10GB+1GB X710 rNDC 10GB X10 rNDC 10GB X10 rNDC 10GB XL710 QSFP DP LP PCIe 40GE XL710 QSFP DP FH PCIe 40GE X550 DP BT PCIe 2 x 10 Gb X550 DP BT LP PCIe 2 x 10 Gb XXV710 Fab A/B Mezz 25 Gb (for MX platforms)
Mellanox	 ConnectX-3 Pro 10G Mezz 10GB ConnectX-4 LX 25GE SFP DP rNDC 25GB ConnectX-4 LX 25GE DP FH PCIe 25GB ConnectX-4 LX 25GE DP LP PCIe 25GB ConnectX-4 LX Fab A/B Mezz 25GB (for MX platforms)
Qlogic	 57810 PCle 10GB 57810 bNDC 10GB 57810 Mezz 10GB 57800 rNDC 10GB+1GB 57840 rNDC 10GB 57840 bNDC 10GB GME2662 Mezz FC16 QLE 2692 SP FC16 Gen 6 HBA FH PCle FC16 SP FC16 Gen 6 HBA LP PCle FC16 QLE 2690 DP FC16 Gen 6 HBA FH PCle FC16 QLE 2690 DP FC16 Gen 6 HBA FH PCle FC16 QLE 2742 DP FC32 Gen 6 HBA FH PCle FC32 DP FC32 Gen 6 HBA LP PCle FC32 QLE2740 PCle FC32 QME2692-DEL Fab C Mezz FC16 (for MX platforms) QL41262HMKR-DE Fab A/B Mezz 25 Gb (for MX platforms) QL41232HMKR-DE Fab A/B Mezz 25 Gb (for MX platforms) QLogic 1x32Gb QLE2770 FC HBA
Emulex	 LPe15002B-M8 (FH) PCIe FC8 LPe15002B-M8 (LP) PCIe FC8 LPe15000B-M8 (FH) PCIe FC8 LPe15000B-M8 (LP) PCIe FC8



Table 43. Supported cards for I/O Identity Optimization (continued)

Manufacturer	Туре
	LPe31000-M6-SP PCIe FC16
	LPe31002-M6-D DP PCIe FC16
	LPe32000-M2-D SP PCIe FC32
	LPe32002-M2-D DP PCIe FC32
	• LPe31002-D Fab C Mezz FC16 (for MX platforms)
	• LPe32002-D Fab C Mezz FC32 (for MX platforms)
	• LPe35002-M2 FC32 2-Port
	• LPe35000-M2 FC32 1-Port

Supported NIC firmware versions for IO Identity Optimization

In 14th generation Dell PowerEdge servers, the required NIC firmware is available by default. The following table provides the NIC firmware versions for the I/O identity optimization feature.

Virtual or Remote assigned Address and Persistence Policy behavior when iDRAC is set to Remote-Assigned Address mode or Console mode

The following table describes the Virtual Address Management (VAM) configuration and Persistence Policy behavior, and the dependencies.

Remote assigned Address Feature State in OME Modular	Mode set in iDRAC	IO Identity Feature State in iDRAC	SCP	Persistence Policy	Clear Persistence Policy — Virtual Address
Remote-Assigned Address enabled	RemoteAssignedAd dress Mode	Enabled	Virtual address management (VAM) configured	Configured VAM persists	Set to Remote assigned Address
Remote-Assigned Address enabled	RemoteAssignedAd dress Mode	Enabled	VAM not configured	Set to Remote assigned Address	No persistence — Is set to Remote assigned Address
Remote-Assigned Address enabled	RemoteAssigned Address Mode	Disabled	Configured using the path provided in Lifecycle Controller	Set to Remote assigned Address for that cycle	No persistence — Is set to Remote assigned Address
Remote-Assigned Address enabled	RemoteAssigned Address Mode	Disabled	VAM not configured	Set to Remote assigned Address	Set to Remote assigned Address
Remote-Assigned Address disabled	RemoteAssigned Address Mode	Enabled	VAM configured	Configured VAM persists	Persistence only — clear is not possible
Remote-Assigned Address disabled	RemoteAssigned Address Mode	Enabled	VAM not configured	Set to hardware MAC address	No persistence supported. Depends on card behavior
Remote-Assigned Address disabled	RemoteAssigned Address Mode	Disabled	Configured using the path provided in Lifecycle Controller	Lifecycle Controller configuration persists for that cycle	No persistence supported. Depends on card behavior
Remote-Assigned Address disabled	RemoteAssigned Address Mode	Disabled	VAM not configured	Set to hardware MAC address	Set to hardware MAC address

Table 44. Virtual/Remote-Assigned Address and Persistence Policy behavior

212 Inventorying, monitoring, and configuring network devices

Inserido ao protocolo **22.951.206-4** por: **Pedro Henrique Golin Linhares** em: 31/03/2025 12:49. A autenticidade deste documento pode ser validada no endereço: https://www.eprotocolo.pr.gov.br/spiweb/validarDocumento com o código: fc1c44dfe0686d5547ffd691199be699.



Table 44. Virtual/Remote-Assigned Address and Persistence Policy behavior (continued)

Remote assigned Address Feature State in OME Modular	Mode set in iDRAC	IO Identity Feature State in iDRAC	SCP	Persistence Policy	Clear Persistence Policy — Virtual Address
Remote-Assigned Address enabled	Console Mode	Enabled	VAM configured	Configured VAM persists	Both persistence and clear must work
Remote-Assigned Address enabled	Console Mode	Enabled	VAM not configured	Set to hardware MAC address	Set to hardware MAC address
Remote-Assigned Address enabled	Console Mode	Disabled	Configured using the path provided in Lifecycle Controller	Lifecycle Controller configuration persists for that cycle	No persistence supported. Depends on card behavior
Remote-Assigned Address disabled	Console Mode	Enabled	VAM configured	Configured VAM persists	Both persistence and clear must work
Remote-Assigned Address disabled	Console Mode	Enabled	VAM not configured	Set to hardware MAC address	Set to hardware MAC address
Remote-Assigned Address disabled	Console Mode	Disabled	Configured using the path provided in Lifecycle Controller	Lifecycle Controller configuration persists for that cycle	No persistence supported. Depends on card behavior
Remote-Assigned Address enabled	Console Mode	Disabled	VAM not configured	Set to hardware MAC address	Set to hardware MAC address

System behavior for FlexAddress and IO Identity

Table 45. System behavior for FlexAddress and I/O Identity

Туре	FlexAddress Feature State in CMC	IO Identity Feature State in iDRAC	Availability of Remote Agent VA for the Reboot Cycle	VA Programming Source	Reboot Cycle VA Persistence Behavior
Server with FA-equivalent	Enabled	Disabled		FlexAddress from CMC	Per FlexAddress spec
Persistence	N/A, Enabled, or Enabled Disabled		Yes - New or Persisted	Remote Agent Virtual Address	Per FlexAddress spec
			No	Virtual Address Cleared	
	Disabled	Disabled			
Server with VAM Persistence Policy Feature	Enabled	Disabled		FlexAddress from CMC	Per FlexAddress spec
	Enabled Ena	Enabled	Yes — New or Persisted	Remote Agent Virtual Address	Per Remote Agent Policy Setting
			No	FlexAddress from CMC	Per FlexAddress spec
	Disabled Enabled	Yes — New or Persisted	Remote Agent Virtual Address	Per Remote Agent Policy Setting	
			No	Virtual Address Cleared	
	Disabled	Disabled			



Enabling or disabling IO Identity Optimization

Normally, after the system boots, the devices are configured and then after a reboot the devices are initialized. You can enable the I/O Identity Optimization feature to achieve boot optimization. If it is enabled, it sets the virtual address, initiator, and storage target attributes after the device is reset and before it is initialized, thus eliminating a second BIOS restart. The device configuration and boot operation occur in a single system start and is optimized for boot time performance.

Before enabling I/O identity optimization, make sure that:

- You have the Login, Configure, and System Control privileges.
- BIOS, iDRAC, and network cards are updated to the latest firmware.

After enabling I/O Identity Optimization feature, export the Server Configuration Profile file from iDRAC, modify the required I/O Identity attributes in the SCP file, and import the file back to iDRAC.

For the list of I/O Identity Optimization attributes that you can modify in the SCP file, see the *NIC Profile* document available at https://www.dell.com/support.

(i) NOTE: Do not modify non I/O Identity Optimization attributes.

Enabling or disabling IO Identity Optimization using web interface

To enable or disable I/O Identity Optimization:

 In the iDRAC Web interface, go to Configuration > System Settings > Hardware Settings > I/O Identity Optimization.

The I/O Identity Optimization page is displayed.

- 2. Click the I/O Identity Optimization tab, select the Enable option to enable this feature. To disable, clear this option.
- 3. Click Apply to apply the setting.

Enabling or disabling IO Identity Optimization using RACADM

To enable I/O Identity Optimization, use the command:

racadm set idrac.ioidopt.IOIDOptEnable Enabled

After enabling this feature, you must restart the system for the settings to take effect.

To disable I/O Identity Optimization, use the command:

racadm set idrac.ioidopt.IOIDOptEnable Disabled

To view the I/O Identity Optimization setting, use the command:

racadm get iDRAC.IOIDOpt

SSD Wear Threshold

iDRAC provides you the ability to configure thresholds of Remaining Rated Write Endurance for all SSD's and Available Spare of NVMe PCIe SSDs.

When SSD Remaining Rated Write Endurance and NVMe PCIe SSD Available Spare values are less than the threshold, then iDRAC logs this event in the LC log and depending on the alert type selection, iDRAC also performs Email alert, SNMP Trap, IPMI Alert, Logging in Remote Syslog, WS Eventing and OS log.

iDRAC alerts the user when the SSD Remaining Rated Write Endurance goes below the set threshold, so that the system admin can take a backup of SSD or replace it.

For only NVMe PCle SSDs, iDRAC displays **Available Spare** and provide a threshold to warn. The **Available Spare** is not available for SSDs which are connected behind PERC and HBA.



Configuring SSD Wear Threshold alert features using web interface

To configure Remaining Rated Write Endurance and Available Spare Alert Threshold using web interface:

- 1. In the iDRAC Web interface, go to Configuration > System Settings > Hardware Settings > SSD Wear Thresholds. The SSD Wear Thresholds page is displayed.
- Remaining Rated Write Endurance You can set the value between 1-99%. The default value is 10%. Alert type for this feature is SSD Wear Write Endurance and security alert is Warning as a result of threshold event..
- Available Spare Alert Threshold You can set the value between 1-99%. The default value is 10%. Alert type for this feature is SSD Wear Available Spare and security alert is Warning as a result of threshold event.

Configuring SSD Wear Threshold alert features using RACADM

To configure Remaining Rated Write Endurance, use the command:

racadm set System.Storage.RemainingRatedWriteEnduranceAlertThreshold n

, where n= 1 to 99%.

To configure Available Spare Alert Threshold, use the command:

racadm System.Storage.AvailableSpareAlertThreshold n

```
, where n= 1 to 99%.
```

Configuring persistence policy settings

Using IO identity, you can configure policies specifying the system reset and power cycle behaviors that determine the persistence or clearance of the virtual address, initiator, and storage target settings. Each individual persistence policy attribute applies to all ports and partitions of all applicable devices in the system. The device behavior changes between auxiliary powered devices and non-auxiliary powered devices.

() NOTE: The Persistence Policy feature may not work when set to default, if the VirtualAddressManagement attribute is set to FlexAddress (not for MX platforms) or RemoteAssignedAddress (for MX platforms) mode on iDRAC and if the FlexAddress or Remote-Assigned Address feature is disabled in CMC (not for MX platforms) or OME Modular (for MX platforms), ensure that you set the VirtualAddressManagement attribute to Console mode in iDRAC or enable the FlexAddress or Remote-Assigned Address feature in CMC or OME Modular.

You can configure the following persistence policies:

- Virtual Address: Auxiliary powered devices
- Virtual Address: Non-Auxiliary powered devices
- Initiator
- Storage target

Before applying the persistence policy, make sure to:

- Inventory the network hardware at least once, that is, enabled Collect System Inventory On Restart.
- Enable I/O Identity Optimization.

Events are logged to the Lifecycle Controller log when:

- I/O Identity Optimization is enabled or disabled.
- Persistence policy is changed.
- Virtual address, initiator and target values are set based on the policy. A single log entry is logged for the configured devices and the values that are set for those devices when the policy is applied.

Event actions are enabled for SNMP, email, or WS-eventing notifications. Logs are also included in the remote syslogs.



Default values for persistence policy

Table 46. Default values for persistence policy

Persistence Policy	AC Power Loss	Cold Boot	Warm Boot
Virtual Address: Auxiliary Powered Devices	Not selected	Selected	Selected
Virtual Address: Non-Auxiliary Powered Devices	Not selected	Not selected	Selected
Initiator	Selected	Selected	Selected
Storage Target	Selected	Selected	Selected

NOTE: When a persistent policy is disabled and when you perform the action to lose the virtual address, re-enabling the persistent policy does not retrieve the virtual address. You must set the virtual address again after you enable the persistent policy.

() NOTE: If there is a persistence policy in effect and the virtual addresses, initiator, or storage targets are set on a CNA-device partition, do not reset or clear the values configured for virtual addresses, initiator, and storage targets before changing the VirtualizationMode or the personality of the partition. The action is performed automatically when you disable the persistence policy. You can also use a configuration job to explicitly set the virtual address attributes to 0s and the initiator and storage targets values as defined in iSCSI initiator and storage target default values.

Configuring persistence policy settings using iDRAC web interface

To configure the persistence policy:

- In the iDRAC Web interface, go to Configuration > System Settings > Hardware Settings > I/O Identity Optimization.
- 2. Click I/O Identity Optimization tab.
- 3. In the **Persistence Policy** section, select one or more of the following for each persistence policy:
 - Warm Reset The virtual address or target settings persist when warm reset condition occurs.
 - Cold Reset The virtual address or target settings persist when cold reset conditions occur.
 - AC Power Loss The virtual address or target settings persist when AC power loss conditions occur.
- 4. Click Apply.

The persistence policies are configured.

Configuring persistence policy settings using RACADM

To set persistence policy, use the following racadm object with the **set** sub command:

- For virtual addresses, use iDRAC.IOIDOpt.VirtualAddressPersistencePolicyAuxPwrd and iDRAC.IOIDOpt.VirtualAddressPersistencePolicyNonAuxPwrd objects
- For initiator, use iDRAC.IOIDOPT.InitiatorPersistencePolicy object
- For storage targets, use **iDRAC.IOIDOpt.StorageTargetPersistencePolicy** object

For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

iSCSI initiator and storage target default values

The following tables provide the list of default values for iSCSI initiator and storage targets when the persistence policies are cleared.

Table 47. iSCSI initiator —default values

iSCSI Initiator	Default Values in IPv4 mode	Default Values in IPv6 mode
lscsilnitiatorlpAddr	0.0.0.0	



Table 47. iSCSI initiator —default values (continued)

iSCSI Initiator	Default Values in IPv4 mode	Default Values in IPv6 mode	
lscsilnitiatorlpv4Addr	0.0.0.0	0.0.0	
lscsilnitiatorlpv6Addr			
lscsilnitiatorSubnet	0.0.0.0	0.0.0.0	
IscsilnitiatorSubnetPrefix	0	0	
lscsilnitiatorGateway	0.0.0.0		
lscsilnitiatorlpv4Gateway	0.0.0.0	0.0.0.0	
lscsilnitiatorlpv6Gateway			
IscsilnitiatorPrimDns	0.0.0.0	::	
lscsilnitiatorlpv4PrimDns	0.0.0.0	0.0.0.0	
lscsilnitiatorlpv6PrimDns		::	
IscsilnitiatorSecDns	0.0.0.0		
lscsilnitiatorlpv4SecDns	0.0.0.0	0.0.0.0	
lscsilnitiatorlpv6SecDns		::	
IscsilnitiatorName	Value Cleared	Value Cleared	
lscsilnitiatorChapld	Value Cleared	Value Cleared	
lscsilnitiatorChapPwd	Value Cleared	Value Cleared	
IPVer	lpv4	lpv6	

Table 48. ISCSI storage target attributes — default values

iSCSI Storage Target Attributes	Default Values in IPv4 mode	Default Values in IPv6 mode
ConnectFirstTgt	Disabled	Disabled
FirstTgtlpAddress	0.0.0.0	::
FirstTgtTcpPort	3260	3260
FirstTgtBootLun	0	0
FirstTgtlscsiName	Value Cleared	Value Cleared
FirstTgtChapId	Value Cleared	Value Cleared
FirstTgtChapPwd	Value Cleared	Value Cleared
FirstTgtlpVer	lpv4	
ConnectSecondTgt	Disabled	Disabled
SecondTgtlpAddress	0.0.0.0	



Table 48. ISCSI storage target attributes — default values (continued)

iSCSI Storage Target Attributes	Default Values in IPv4 mode	Default Values in IPv6 mode
SecondTgtTcpPort	3260	3260
SecondTgtBootLun	0	0
SecondTgtIscsiName	Value Cleared	Value Cleared
SecondTgtChapId	Value Cleared	Value Cleared
SecondTgtChapPwd	Value Cleared	Value Cleared
SecondTgtlpVer	Ірv4	



Managing storage devices

Starting with iDRAC 3.15.15.15 release, iDRAC supports Boot Optimized Storage Solution (BOSS) controller in the 14th generation of PowerEdge servers. BOSS controllers are designed specifically for booting the operating system of the server. These controllers support limited RAID features and the configuration is staged.

Starting with iDRAC 4.30.30.30 release, iDRAC supports PERC 11, HBA 11, and BOOS 1.5 for AMD systems.

NOTE: BOSS controllers support only RAID level1.

NOTE: For BOSS Controllers, the complete VD information may not be available when both PD's are plugged-out and plugged-in back.

INOTE: PERC 11 and later controllers support Hardware Root of Trust (RoT).

iDRAC has expanded its agent-free management to include direct configuration of the PERC controllers. It enables you to remotely configure the storage components attached to your system at run-time. These components include RAID and non-RAID controllers and the channels, ports, enclosures, and disks attached to them. For the PowerEdge Rx4xx/Cx4xx servers, PERC 9 and PERC 10 controllers are supported. For PowerEdge Rx5xx/Cx5xx AMD platform servers, PERC 11 is supported.

The complete storage subsystem discovery, topology, health monitoring, and configuration are accomplished in the Comprehensive Embedded Management (CEM) framework by interfacing with the internal and external PERC controllers through the MCTP protocol over I2C interface. For real-time configuration, CEM supports PERC9 controllers and above. The firmware version for PERC9 controllers must be 9.1 or later.

NOTE: The Software RAID (SWRAID) is not supported by CEM and thus is not supported in the iDRAC GUI. SWRAID can be managed using either RACADM, WSMAN or Redfish.

Using iDRAC, you can perform most of the functions that are available in OpenManage Storage Management including real-time (no reboot) configuration commands (for example, create virtual disk). You can completely configure RAID before installing the operating system.

You can configure and manage the controller functions without accessing the BIOS. These functions include configuring virtual disks and applying RAID levels and hot spares for data protection. You can initiate many other controller functions such as rebuilds and troubleshooting. You can protect your data by configuring data-redundancy or assigning hot spares.

The storage devices are:

- Controllers Most operating systems do not read and write data directly from the disks, but instead send read and write instructions to a controller. The controller is the hardware in your system that interacts directly with the disks to write and retrieve data. A controller has connectors (channels or ports) which are attached to one or more physical disks or an enclosure containing physical disks. RAID controllers can span the boundaries of the disks to create an extended amount of storage space— or a virtual disk using the capacity of more than one disk. Controllers also perform other tasks, such as initiating rebuilds, initializing disks, and more. To complete their tasks, controllers require special software known as firmware and drivers. In order to function properly, the controller must have the minimum required version of the firmware and the drivers installed. Different controllers have different characteristics in the way they read and write data and execute tasks. It is helpful to understand these features to most efficiently manage the storage.
- Physical disks or physical devices Reside within an enclosure or are attached to the controller. On a RAID controller, physical disks or devices are used to create virtual disks.
- Virtual disk It is storage created by a RAID controller from one or more physical disks. Although a virtual disk may be created from several physical disks, it is viewed by the operating system as a single disk. Depending on the RAID level used, the virtual disk may retain redundant data if there is a disk failure or have particular performance attributes. Virtual disks can only be created on a RAID controller.
- Enclosure It is attached to the system externally while the backplane and its physical disks are internal.
- Backplane It is similar to an enclosure. In a Backplane, the controller connector and physical disks are attached to the enclosure, but it does not have the management features (temperature probes, alarms, and so on) associated with external enclosures. Physical disks can be contained in an enclosure or attached to the backplane of a system.

() NOTE: In any MX chassis which contains storage sleds and compute sleds, iDRAC pertaining to any of the compute sleds in that chassis will report all storage sleds (both assigned and unassigned). If any one of the assigned or unassigned blades are in Warning or Critical heath state, the blade controller also reports the same status.



In addition to managing the physical disks contained in the enclosure, you can monitor the status of the fans, power supply, and temperature probes in an enclosure. You can hot-plug enclosures. Hot-plugging is defined as adding of a component to a system while the operating system is still running.

The physical devices connected to the controller must have the latest firmware. For the latest supported firmware, contact your service provider.

Storage events from PERC are mapped to SNMP traps and WSMan events as applicable. Any changes to the storage configurations are logged in the Lifecycle Log.

Table 49. PERC capability

PERC Capability	CEM configuration Capable Controller (PERC 9.1 or later)	CEM configuration Non-capable Controller (PERC 9.0 and lower)
Real-time	() NOTE: For PowerEdge Rx5xx/ Cx5xx servers, PERC 9, PERC 10, and PERC 11 controllers are supported.	Configuration is applied. An error message is displayed. Job creation is not successful and you cannot create real- time jobs using Web interface.
	If there is no existing pending or scheduled jobs for the controller, then configuration is applied.	
	If there are pending or scheduled jobs for that controller, then the jobs have to be cleared or you must wait for the jobs to be completed before applying the configuration at run-time. Run-time or real-time means, a reboot is not required.	
Staged	If all the set operations are staged, the configuration is staged and applied after reboot or it is applied at real-time.	Configuration is applied after reboot

Topics:

- Understanding RAID concepts
- Supported controllers
- Supported enclosures
- Summary of supported features for storage devices
- Inventorying and monitoring storage devices
- Viewing storage device topology
- Managing physical disks
- Managing virtual disks
- RAID Configuration Features
- Managing controllers
- Managing PCIe SSDs
- Managing enclosures or backplanes
- Choosing operation mode to apply settings
- Viewing and applying pending operations
- Storage devices apply operation scenarios
- Blinking or unblinking component LEDs
- Warm reboot

Understanding RAID concepts

Storage Management uses the Redundant Array of Independent Disks (RAID) technology to provide Storage Management capability. Understanding Storage Management requires an understanding of RAID concepts, as well as some familiarity with how the RAID controllers and operating system view disk space on your system.

220 Managing storage devices



What is RAID

RAID is a technology for managing the storage of data on the physical disks that reside or are attached to the system. A key aspect of RAID is the ability to span physical disks so that the combined storage capacity of multiple physical disks can be treated as a single, extended disk space. Another key aspect of RAID is the ability to maintain redundant data which can be used to restore data in the event of a disk failure. RAID uses different techniques, such as striping, mirroring, and parity, to store and reconstruct data. There are different RAID levels that use different methods for storing and reconstructing data. The RAID levels have different characteristics in terms of read/write performance, data protection, and storage capacity. Not all RAID levels maintain redundant data, which means for some RAID levels lost data cannot be restored. The RAID level you choose depends on whether your priority is performance, protection, or storage capacity.

() NOTE: The RAID Advisory Board (RAB) defines the specifications used to implement RAID. Although RAB defines the RAID levels, commercial implementation of RAID levels by different vendors may vary from the actual RAID specifications. An implementation of a particular vendor may affect the read and write performance and the degree of data redundancy.

Hardware and software RAID

RAID can be implemented with either hardware or software. A system using hardware RAID has a RAID controller that implements the RAID levels and processes data reads and writes to the physical disks. When using software RAID provided by the operating system, the operating system implements the RAID levels. For this reason, using software RAID by itself can slow the system performance. You can, however, use software RAID along with hardware RAID volumes to provide better performance and variety in the configuration of RAID volumes. For example, you can mirror a pair of hardware RAID 5 volumes across two RAID controllers to provide RAID controller redundancy.

RAID concepts

RAID uses particular techniques for writing data to disks. These techniques enable RAID to provide data redundancy or better performance. These techniques include:

- Mirroring Duplicating data from one physical disk to another physical disk. Mirroring provides data redundancy by maintaining two copies of the same data on different physical disks. If one of the disks in the mirror fails, the system can continue to operate using the unaffected disk. Both sides of the mirror contain the same data always. Either side of the mirror can act as the operational side. A mirrored RAID disk group is comparable in performance to a RAID 5 disk group in read operations but faster in write operations.
- Striping Disk striping writes data across all physical disks in a virtual disk. Each stripe consists of consecutive virtual disk data addresses that are mapped in fixed-size units to each physical disk in the virtual disk using a sequential pattern. For example, if the virtual disk includes five physical disks, the stripe writes data to physical disks one through five without repeating any of the physical disks. The amount of space consumed by a stripe is the same on each physical disk. The portion of a stripe that resides on a physical disk is a stripe element. Striping by itself does not provide data redundancy. Striping in combination with parity does provide data redundancy.
- Stripe size The total disk space consumed by a stripe not including a parity disk. For example, consider a stripe that contains 64KB of disk space and has 16KB of data residing on each disk in the stripe. In this case, the stripe size is 64KB and the stripe element size is 16KB.
- Stripe element A stripe element is the portion of a stripe that resides on a single physical disk.
- Stripe element size The amount of disk space consumed by a stripe element. For example, consider a stripe that contains 64KB of disk space and has 16KB of data residing on each disk in the stripe. In this case, the stripe element size is 16KB and the stripe size is 64KB.
- Parity Parity refers to redundant data that is maintained using an algorithm in combination with striping. When one of the striped disks fails, the data can be reconstructed from the parity information using the algorithm.
- Span A span is a RAID technique used to combine storage space from groups of physical disks into a RAID 10, 50, or 60 virtual disk.

RAID levels

Each RAID level uses some combination of mirroring, striping, and parity to provide data redundancy or improved read and write performance. For specific information on each RAID level, see Choosing raid levels.



Organizing data storage for availability and performance

RAID provides different methods or RAID levels for organizing the disk storage. Some RAID levels maintain redundant data so that you can restore data after a disk failure. Different RAID levels also entail an increase or decrease in the I/O (read and write) performance of a system.

Maintaining redundant data requires the use of additional physical disks. The possibility of a disk failure increases with an increase in the number of disks. Since the differences in I/O performance and redundancy, one RAID level may be more appropriate than another based on the applications in the operating environment and the nature of the data being stored.

When choosing a RAID level, the following performance and cost considerations apply:

- Availability or fault-tolerance Availability or fault-tolerance refers to the ability of a system to maintain operations and provide access to data even when one of its components has failed. In RAID volumes, availability or fault-tolerance is achieved by maintaining redundant data. Redundant data includes mirrors (duplicate data) and parity information (reconstructing data using an algorithm).
- Performance Read and write performance can be increased or decreased depending on the RAID level you choose. Some RAID levels may be more appropriate for particular applications.
- Cost efficiency Maintaining the redundant data or parity information associated with RAID volumes requires additional disk space. In situations where the data is temporary, easily reproduced, or non-essential, the increased cost of data redundancy may not be justified.
- Mean Time Between Failure (MTBF) Using additional disks to maintain data redundancy also increases the chance of disk failure at any given moment. Although this option cannot be avoided in situations where redundant data is a requirement, it does have implications on the workload of the system support staff within your organization.
- Volume Volume refers to a single disk non-RAID virtual disk. You can create volumes using external utilities like the O-ROM <Ctrl> <r>. Storage Management does not support the creation of volumes. However, you can view volumes and use drives from these volumes for creation of new virtual disks or Online Capacity Expansion (OCE) of existing virtual disks, provided free space is available.

Choosing RAID levels

You can use RAID to control data storage on multiple disks. Each RAID level or concatenation has different performance and data protection characteristics.

(i) NOTE: The H3xx PERC controllers do not support RAID levels 6 and 60.

The following topics provide specific information on how each RAID level store data as well as their performance and protection characteristics:

- Raid level 0 (striping)
- Raid level 1 (mirroring)
- Raid level 5 (striping with distributed parity)
- Raid level 6 (striping with additional distributed parity)
- Raid level 50 (striping over raid 5 sets)
- Raid level 60 (striping over raid 6 sets)
- Raid level 10 (striping over mirror sets)

RAID level 0 - striping

RAID 0 uses data striping, which is writing data in equal-sized segments across the physical disks. RAID 0 does not provide data redundancy.





RAID 0 characteristics:

- Groups *n* disks as one large virtual disk with a capacity of (smallest disk size) **n* disks.
- Data is stored to the disks alternately.
- No redundant data is stored. When a disk fails, the large virtual disk fails with no means of rebuilding the data.
- Better read and write performance.

RAID level 1 - mirroring

RAID 1 is the simplest form of maintaining redundant data. In RAID 1, data is mirrored or duplicated on one or more physical disks. If a physical disk fails, data can be rebuilt using the data from the other side of the mirror.



RAID 1 characteristics:

- Groups *n* + *n* disks as one virtual disk with the capacity of *n* disks. The controllers currently supported by Storage Management allow the selection of two disks when creating a RAID 1. Because these disks are mirrored, the total storage capacity is equal to one disk.
- Data is replicated on both the disks.
- When a disk fails, the virtual disk still works. The data is read from the mirror of the failed disk.
- Better read performance, but slightly slower write performance.



- Redundancy for protection of data.
- RAID 1 is more expensive in terms of disk space since twice the number of disks are used than required to store the data without redundancy.

RAID level 5 or striping with distributed parity

RAID 5 provides data redundancy by using data striping in combination with parity information. Rather than dedicating a physical disk to parity, the parity information is striped across all physical disks in the disk group.



RAID 5 characteristics:

- Groups *n* disks as one large virtual disk with a capacity of (*n*-1) disks.
- Redundant information (parity) is alternately stored on all disks.
- When a disk fails, the virtual disk still works, but it is operating in a degraded state. The data is reconstructed from the surviving disks.
- Better read performance, but slower write performance.
- Redundancy for protection of data.

RAID level 6-striping with additional distributed parity

RAID 6 provides data redundancy by using data striping in combination with parity information. Similar to RAID 5, the parity is distributed within each stripe. RAID 6, however, uses an additional physical disk to maintain parity, such that each stripe in the disk group maintains two disk blocks with parity information. The additional parity provides data protection in the event of two disk failures. In the following image, the two sets of parity information are identified as **P** and **Q**.





RAID 6 characteristics:

- Groups *n* disks as one large virtual disk with a capacity of (*n*-2) disks.
- Redundant information (parity) is alternately stored on all disks.
- The virtual disk remains functional with up to two disk failures. The data is reconstructed from the surviving disks.
- Better read performance, but slower write performance.
- Increased redundancy for protection of data.
- Two disks per span are required for parity. RAID 6 is more expensive in terms of disk space.

RAID level 50 - striping over RAID 5 sets

RAID 50 is striping over more than one span of physical disks. For example, a RAID 5 disk group that is implemented with three physical disks and then continues on with a disk group of three more physical disks would be a RAID 50.

It is possible to implement RAID 50 even when the hardware does not directly support it. In this case, you can implement more than one RAID 5 virtual disks and then convert the RAID 5 disks to dynamic disks. You can then create a dynamic volume that is spanned across all RAID 5 virtual disks.





RAID 50 characteristics:

- Groups n^*s disks as one large virtual disk with a capacity of $s^*(n-1)$ disks, where s is the number of spans and n is the number of disks within each span.
- Redundant information (parity) is alternately stored on all disks of each RAID 5 span.
- Better read performance, but slower write performance.
- Requires as much parity information as standard RAID 5.
- Data is striped across all spans. RAID 50 is more expensive in terms of disk space.

RAID level 60 - striping over RAID 6 sets

RAID 60 is striping over more than one span of physical disks that are configured as a RAID 6. For example, a RAID 6 disk group that is implemented with four physical disks and then continues on with a disk group of four more physical disks would be a RAID 60.





RAID 60 characteristics:

- Groups n^*s disks as one large virtual disk with a capacity of $s^*(n-2)$ disks, where s is the number of spans and n is the number of disks within each span.
- Redundant information (parity) is alternately stored on all disks of each RAID 6 span.
- Better read performance, but slower write performance.
- Increased redundancy provides greater data protection than a RAID 50.
- Requires proportionally as much parity information as RAID 6.
- Two disks per span are required for parity. RAID 60 is more expensive in terms of disk space.

RAID level 10 - striped with mirrors

The RAB considers RAID level 10 to be an implementation of RAID level 1. RAID 10 combines mirrored physical disks (RAID 1) with data striping (RAID 0). With RAID 10, data is striped across multiple physical disks. The striped disk group is then mirrored onto another set of physical disks. RAID 10 can be considered a *mirror of stripes*.





RAID 10 characteristics:

- Groups *n* disks as one large virtual disk with a capacity of (n/2) disks, where *n* is an even integer.
- Mirror images of the data are striped across sets of physical disks. This level provides redundancy through mirroring.
- When a disk fails, the virtual disk still works. The data is read from the surviving mirrored disk.
- Improved read performance and write performance.
- Redundancy for protection of data.

Comparing RAID level performance

The following table compares the performance characteristics associated with the more common RAID levels. This table provides general guidelines for choosing a RAID level. Evaluate your specific environment requirements before choosing a RAID level.

RAID Level	Data Redundancy	Read Performance	Write Performance	Rebuild Performance	Minimum Disks Required	Suggested Uses
RAID 0	None	Very Good	Very Good	N/A	Ν	Noncritical data.
RAID 1	Excellent	Very Good	Good	Good	2N (N = 1)	Small databases, database logs, and critical information.
RAID 5	Good	Sequential reads: good. Transactional reads: Very good	Fair, unless using writeback cache	Fair	N + 1 (N = at least two disks)	Databases and other read intensive transactional uses.
RAID 10	Excellent	Very Good	Fair	Good	2N x X	Data intensive environments (large records).

Table 50. RAID level performance comparison



Table 50. RAID level performance comparison (continued)

RAID Level	Data Redundancy	Read Performance	Write Performance	Rebuild Performance	Minimum Disks Required	Suggested Uses			
RAID 50	Good	Very Good	Fair	Fair	N + 2 (N = at least 4)	Medium sized transactional or data intensive uses.			
RAID 6	Excellent	Sequential reads: good. Transactional reads: Very good	Fair, unless using writeback cache	Poor	N + 2 (N = at least two disks)	Critical information. Databases and other read intensive transactional uses.			
RAID 60	Excellent	Very Good	Fair	Poor	X x (N + 2) (N = at least 2)	Critical information. Medium sized transactional or data intensive uses.			
N = Number of physical disks									
X = Number of R	AID sets								

Supported controllers

Supported RAID controllers

The iDRAC interfaces support the following BOSS controllers:

- BOSS-S1 Adapter
- BOSS-S1 Modular (for blade servers)
- BOSS-S2 Adapter

The iDRAC interfaces support the following PERC11 controllers:

- PERC H755 Adapter
- PERC H755 Front
- PERC H755N Front
- The iDRAC interfaces support the following PERC10 controllers:
- PERC H740P Mini
- PERC H740P Adapter
- PERC H840 Adapter
- PERC H745P MX

The iDRAC interfaces support the following PERC9 controllers:

- PERC H330 Mini
- PERC H330 Adapter
- PERC H730P Mini
- PERC H730P Adapter
- PERC H730P MX

Supported non-RAID controllers

The iDRAC interface supports 12 Gbps SAS HBA external controller and HBA330 Mini or Adapter controllers.

iDRAC supports HBA330 MMZ, HBA330 MX adapters.



Supported enclosures

iDRAC supports MD1400 and MD1420 enclosures.

(i) NOTE: Redundant Array of Inexpensive Disks (RBODS) that are connected to HBA controllers are not supported.

(i) NOTE: PERC H480 with version 10.1 or greater, firmware supports up to 4 enclosures per port.

Summary of supported features for storage devices

The following tables provide the features supported by the storage devices through iDRAC.

Table 51. Supported features for storage controllers

Feature	PERC 11			PERC 10 P			PERC 9				
	H755 Front	H755N Front	H755 Adapter	H740P Mini	H740P Adapter	H840 Adapter	H330 Mini	H330 Adapter	H730P Mini	H730P Adapter	FD33xS
Assign or unassign physical disk as a global hot spare	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time
Convert to RAID	Not applicabl e	Not applicabl e	Not applicabl e	Not applicabl e	Not applicabl e	Not applicabl e	Not applicabl e	Not applicabl e	Not applicabl e	Not applicabl e	Not applicabl e
Convert to RAID/ Non RAID,	Real- time (convert s drive to nonRAID ePD-PT volume)	Real- time (convert s drive to nonRAID ePD-PT volume)	Real- time (convert s drive to nonRAID ePD-PT volume)	Real- time (only supporte d in eHBA controlle r mode, converts drive to nonRAID ePD-PT volume)	Real- time (only supporte d in eHBA controlle r mode, converts drive to nonRAID ePD-PT volume)	Real- time (only supporte d in eHBA controlle r mode, converts drive to nonRAID ePD-PT volume)	Real- time	Real- time	Real- time	Real- time	Real- time
Rebuild	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time
Cancel Rebuild	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time
Create virtual disks	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time
Rename virtual disks	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time
Edit virtual disks cache policies	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time



Table 51. Supported features for storage controllers (continued)

Feature	PERC 11			PERC 10			PERC 9				
	H755 Front	H755N Front	H755 Adapter	H740P Mini	H740P Adapter	H840 Adapter	H330 Mini	H330 Adapter	H730P Mini	H730P Adapter	FD33xS
Check virtual disk consiste ncy	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time
Cancel check consiste ncy	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time
lnitialize virtual disks	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time
Cancel initializat ion	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time
Encrypt virtual disks	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Not applicabl e	Not applicabl e	Real- time	Real- time	Real- time
Assign or unassign dedicate d hot spare	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time
Delete virtual disks	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time
Cancel Backgro und Initializat ion	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time
Online Capacity Expansio n	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time
RAID Level Migratio n	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time
Discard Preserve d Cache	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Not applicabl e	Not applicabl e	Real- time	Real- time	Real- time
Set Patrol Read Mode	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time
Manual Patrol Read Mode	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time



Table 51. Supported features for storage controllers (continued)

Feature	PERC 11			PERC 10			PERC 9				
	H755 Front	H755N Front	H755 Adapter	H740P Mini	H740P Adapter	H840 Adapter	H330 Mini	H330 Adapter	H730P Mini	H730P Adapter	FD33xS
Patrol Read Unconfig ured Areas	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time (only in web interface)				
Check Consiste ncy Mode	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time
Copybac k Mode	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time
Load Balance Mode	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time
Check Consiste ncy Rate	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time
Rebuild Rate	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time
BGI Rate	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time
Reconst ruct Rate	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time
Import foreign configur ation	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time
Auto- import foreign configur ation	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time
Clear foreign configur ation	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time
Reset controlle r configur ation	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time
Create or change security keys	Real- time	Real- time	Real- time	Real- time	Real- time	Real- time	Not applicabl e	Not applicabl e	Real- time	Real- time	Real- time

Inserido ao protocolo **22.951.206-4** por: **Pedro Henrique Golin Linhares** em: 31/03/2025 12:49. A autenticidade deste documento pode ser validada no endereço: https://www.eprotocolo.pr.gov.br/spiweb/validarDocumento com o código: fc1c44dfe0686d5547ffd691199be699.



Table 51. Supported features for storage controllers (continued)

Feature	PERC 11			PERC 10			PERC 9				
	H755 Front	H755N Front	H755 Adapter	H740P Mini	H740P Adapter	H840 Adapter	H330 Mini	H330 Adapter	H730P Mini	H730P Adapter	FD33xS
Secure Enterpri se Key Manger	Staged	Staged	Staged	Staged	Staged	Staged	Not applicabl e	Not applicabl e	Not applicabl e	Not applicabl e	Not applicabl e
Inventor y and remotely monitor the health of PCle SSD devices	Not applicabl e										
Prepare the PCle SSD to be removed	Not applicabl e										
Securely erase the data for PCle SSD	Not applicabl e	Real- time	Not applicabl e								
Configur e Backplan e mode (split/ unified)	Real- time										
Blink or unblink compon ent LEDs	Real- time										
Switch controlle r mode	Not applicabl e	Not applicabl e	Not applicabl e	Staged							
T10PI support for Virtual Disks	Not applicabl e										

(i) NOTE: Added support for

• eHBA mode for PERC 10.2 or greater firmware which supports convert to Non-RAID disks

- convert controller to HBA mode
- RAID 10 uneven span



Table 52. Supported features of storage controllers for MX platforms

Features	PERC 11	PERC 10	PERC 9	
	H755 MX	H745P MX	H730P MX	
Initialize Virtual Disks	Real-time	Real-time	Real-time	
Cancel Initialization	Real-time	Real-time	Real-time	
Encrypt Virtual Disks	Real-time	Real-time	Real-time	
Assign or unassign dedicated hot spare	Real-time	Real-time	Real-time	
Delete Virtual Disks	Real-time	Real-time	Real-time	
Cancel Background Initialization	Real-time	Real-time	Real-time	
Online Capacity Expansion	Real-time	Real-time	Real-time	
RAID Level Migration	Real-time	Real-time	Real-time	
Discard Preserved Cache	Real-time	Real-time	Real-time	
Set Patrol Read Mode	Real-time	Real-time	Real-time	
Manual Patrol Read Mode	Real-time	Real-time	Real-time	
Patrol Read Unconfigured Areas	Real-time	Real-time	Real-time (only in web interface)	
Check Consistency Mode	Real-time	Real-time	Real-time	
Copyback Mode	Real-time	Real-time	Real-time	
Load Balance Mode	Real-time	Real-time	Real-time	
Check Consistency Rate	Real-time	Real-time	Real-time	
Rebuild Rate	Real-time	Real-time	Real-time	
BGI Rate	Real-time	Real-time	Real-time	
Reconstruct Rate	Real-time	Real-time	Real-time	
Import Foreign Configuration	Real-time	Real-time	Real-time	
Auto-import Foreign Configuration	Real-time	Real-time	Real-time	
Clear Foreign Configuration	Real-time	Real-time	Real-time	
Reset Controller Configuration	Real-time	Real-time	Real-time	
Create or change security keys	Real-time	Real-time	Real-time	
Inventory and remotely monitor the health of PCIe SSD devices	Real-time	Not applicable	Not applicable	
Prepare the PCIe SSD to be removed	Not applicable	Not applicable	Not applicable	
Securely erase the data for PCIe SSD	Real-time	Not applicable	Not applicable	
Configure Backplane mode (split/ unified)	Real-time	Not applicable	Not applicable	
Blink or unblink component LEDs	Real-time	Real-time	Real-time	



Table 52. Supported features of storage controllers for MX platforms (continued)

Features	PERC 11	PERC 10	PERC 9	
	H755 MX	H745P MX	Н730Р МХ	
Switch Controller Mode	Not applicable	Not applicable	Staged	
T10PI Support for Virtual Disks	Not applicable	Not applicable	Not applicable	

(i) NOTE: H745P MX supports eHBA mode with PERC 10.2 and higher.

Table 53. Supported features for storage devices

Feature	PCIe SSD	BOSS S1	BOSS S2
Create Virtual Disks	Not applicable	Staged	Staged
Reset Controller Configuration	Not applicable	Staged	Staged
Fast Initialization	Not applicable	Staged	Staged
Delete Virtual Disks	Not applicable	Staged	Staged
Full Initialization	Not applicable	Not applicable	Not applicable
Inventory and remotely monitor the health of PCle SSD devices	Real-time	Not applicable	Not applicable
Prepare the PCIe SSD to be removed	Real-time	Not applicable	Not applicable
Securely erase the data for PCIe SSD	Staged	Not applicable	Not applicable
Blink or unblink component LEDs	Real-time	Not applicable	Real-time
Hot plugging of drives	Real-time	Not applicable	Real-time

Inventorying and monitoring storage devices

You can remotely monitor the health and view the inventory of the following Comprehensive Embedded Management (CEM) enabled storage devices in the managed system using iDRAC web interface:

- RAID controllers, non-RAID controllers, BOSS controllers and PCIe extenders
- Enclosures that include Enclosure Management Modules (EMMs), power supply, fan probe, and temperature probe
- Physical disks
- Virtual disks
- Batteries

The recent storage events and topology of storage devices are also displayed.

Alerts and SNMP traps are generated for storage events. The events are logged in the Lifecycle Log.

() NOTE:

- If you enumerate the enclosure view's WSMan command on a system while one PSU-cable is removed, the primary status of the enclosure view is reported as **Healthy** instead of **Warning**.
- For an accurate inventory of BOSS controllers, ensure that Collect System Inventory On Reboot Operation (CSIOR) is completed. CSIOR is enabled by default.
- The storage health rollup follows the same convention of Dell EMC OpenManage product. For more information see the *OpenManage Server Administrator User's Guide* available at https://www.dell.com/openmanagemanuals.



- Physical disks in system with multiple backplanes may be listed under a different backplane. Use the blink function to identify the disks.
- FQDD of certain Backplanes may not be same in Software Inventory and Hardware Inventory.
- Lifecycle log for PERC controller is not available when the past PERC controller events are being processed and this does not affect the functionality. Past event processing can vary depending on the configuration

Monitoring storage devices using web interface

To view the storage device information using web interface:

- Go to **Storage** > **Overview** > **Summary** to view the summary of the storage components and the recently logged events. This page is automatically refreshed every 30 seconds.
- Go to Storage > Overview > Controllers to view the RAID controller information. The Controllers page is displayed.
- Go to Storage > Overview > Physical Disks to view physical disk information. The Physical Disks page is displayed.
- Go to Storage > Overview > Virtual Disks to view virtual disk information. The Virtual Disks page is displayed.
- Go to Storage > Overview > Enclosures to view the enclosure information. The Enclosures page is displayed.

You can also use filters to view specific device information.

() NOTE:

- The storage hardware list is not displayed in case the system does not have storage devices with CEM support.
- Behavior of non-Dell certified or 3rd party NVMe devices may not be consistent in iDRAC.
- If the NVMe SSDs in the backplane slots support NVMe-MI commands and the I2C connection to backplane slots are fine, the iDRAC discovers these NVMe SSDs and reports them in the interfaces irrespective of the PCI connections to the respective backplane slots.

(i) NOTE:

Туре	Web GUI Support	Other Interfaces support
SATA	Not available	Inventory and RAID configuration
NVMe	Physical disk inventory only	Inventory and RAID configuration

For more information about the displayed properties and to use the filter options, see the iDRAC Online Help.

Monitoring storage devices using RACADM

To view the storage device information, use the storage command.

For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Monitoring backplane using iDRAC settings utility

In the iDRAC Settings utility, go to **System Summary**. The **iDRAC Settings.System Summary** page is displayed. The **Backplane Inventory** section displays the backplane information. For information about the fields, see the *iDRAC Settings Utility Online Help*.

Viewing storage device topology

You can view the hierarchical physical containment view of the key storage components, that is, a list of controllers, enclosures connected to the controller and a link to the physical disk contained in each enclosure. The physical disks attached directly to the controller are also displayed.

To view the storage device topology, go to **Storage** > **Overview**. The **Overview** page displays the hierarchical representation of the storage components in the system. The available options are:

- Controllers
- Physical Disks


- Virtual Disks
- Enclosures

Click the links to view the respective component details.

Managing physical disks

You can perform the following for physical disks:

- View physical disk properties.
- Assign or unassign physical disk as a global hot-spare.
- Convert to RAID capable disk.
- Convert to non-RAID disk.
- Blink or unblink the LED.
- Rebuild physical disk
- Cancel rebuild physical disk
- Cryptographic erase

Assigning or unassigning physical disk as global hot spare

A global hot spare is an unused backup disk that is part of the disk group. Hot spares remain in standby mode. When a physical disk that is used in a virtual disk fails, the assigned hot spare is activated to replace the failed physical disk without interrupting the system or requiring your intervention. When a hot spare is activated, it rebuilds the data for all redundant virtual disks that were using the failed physical disk.

(i) NOTE: From iDRAC v3.00.00.00 or later, you can add global hot spares when virtual disks are not created.

You can change the hot spare assignment by unassigning a disk and choosing another disk as needed. You can also assign more than one physical disk as a global hot spare.

Global hot spares must be assigned and unassigned manually. They are not assigned to specific virtual disks. If you want to assign a hot spare to a virtual disk (it replaces any physical disk that fails in the virtual disk), then see Assigning or unassigning dedicated hot spares.

When deleting virtual disks, all assigned global hot spares may be automatically unassigned when the last virtual disk associated with the controller is deleted.

If you reset the configuration, the virtual disks are deleted and all the hot spares are unassigned.

You must be familiar with the size requirements and other considerations associated with hot spares.

Before assigning a physical disk as a global hot spare:

- Make sure that Lifecycle Controller is enabled.
- If there are no disk drives available in ready state, insert additional disk drives and make sure that the drives are in ready state.
- If physical disks are in non-RAID mode convert them to RAID mode using iDRAC interfaces such as iDRAC web interface, RACADM, Redfish or WSMan, or <CTRL+R>.
 - **NOTE:** During POST, press F2 to enter System Setup or Device Setup. CTRL+R option is no longer supported for PERC 10. CTRL+R only works with PERC 9 while boot mode is set to BIOS.

If you have assigned a physical disk as a global hot spare in Add to Pending Operation mode, the pending operation is created but a job is not created. Then, if you try to unassign the same disk as global hot spare, the assign global hot spare pending operation is cleared.

If you have unassigned a physical disk as a global hot spare in Add to Pending Operation mode, the pending operation is created but a job is not created. Then, if you try to assign the same disk as a global hot spare, the unassign global hot spare pending operation is cleared.

If the last VD is deleted, the global hot spares also returns to ready state.

If a PD is already a global hot spares, user can still assign it again as a global hot spares.



Assigning or unassigning global hot spare using web interface

To assign or unassign a global hot spare for a physical disk drive:

- 1. In the iDRAC web interface, go to **Configuration** > **Storage Configuration**.
- The Storage Configuration page is displayed.
- 2. From the **Controller** drop-down menu, select the controller to view the associated physical disks.
- **3.** Click **Physical Disk Configuration**. All the physical disks associated to the controller are displayed.
- 4. To assign as a global hotspare, from the drop-down menus in the **Action** column, select **Assign Global Hotspare** for one or more physical disks.
- 5. To unassign a hot spare, from the drop-down menus in the **Action** column, select **Unassign Hotspare** for one or more physical disks.
- 6. Click Apply Now.

Depending on your requirement, you can also choose to apply **At Next Reboot** or **At Scheduled Time**. Based on the selected operation mode, the settings are applied.

Assigning or unassigning global hot spare using RACADM

Use the storage command and specify the type as global hot spare.

For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Converting a physical disk to RAID or non-RAID mode

Converting a physical disk to RAID mode enables the disk for all RAID operations. When a disk is in a non-RAID mode, the disk is exposed to the operating system unlike unconfigured good disks and is used in a direct pass-through mode.

PERC 10 is not supported to convert drives to non-RAID. But it is supported in PERC 10.2 and higher versions.

You can convert the physical disk drives to RAID or non-RAID mode by:

- Using iDRAC interfaces such as iDRAC web interface, RACADM, Redfish or WSMan.
- Pressing <Ctrl+R> while restarting the server and selecting the required controller.
- () NOTE: If the physical drives connected to a PERC controller are in non-RAID mode, the size of the disk displayed in the iDRAC interfaces, such as iDRAC GUI, RACADM, Redfish and WSMan, may be slightly less than the actual size of the disk. However, you can use the full capacity of the disk to deploy operating systems.

(i) NOTE:

- Hot plugged disks in PERC H330 are always in non-RAID mode. In other RAID controllers, they are always in RAID mode.
- Hot plugged disks in PERC 11 are either ready or EPD-PT depending on the current auto configure behavior setting.

Converting physical disks to RAID capable or non-RAID mode using the iDRAC web interface

To convert the physical disks to RAID mode or non-RAID mode, perform the following steps:

- 1. In the iDRAC web interface, click Storage > Overview > Physical Disks.
- Click Filter options. Two options are displayed Clear All Filters and Advanced Filter. Click Advanced Filter option. An elaborated list is displayed that allows you to configure different parameters.
- **3.** From the **Group By** drop-down menu, select an enclosure or virtual disks. The parameters associated with the enclosure or the VD are displayed.
- **4.** Click **Apply**, once you select all the desired parameters. For more information about the fields, see the *iDRAC Online Help*. The settings are applied based on the option selected in the operation mode.

Converting physical disks to RAID capable or non-RAID mode using RACADM

Depending on whether you want to convert to RAID or Non-RAID mode, use the following RACADM commands



- To convert to RAID mode, use the racadm storage converttoraid command.
- To convert to Non-RAID mode, use the racadm storage convertiononraid command.

NOTE: On the S140 controller, you can only use the RACADM interface to convert the drives from non-RAID to RAID mode. The supported Software RAID modes are Windows or Linux Mode.

For more information about the commands, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Erasing physical disks

The System Erase feature allows you to erase the contents of the physical drives. This feature is accessible using RACADM or the LC GUI. Physical drives on the server are grouped into two categories.

- Secure erase drives— Includes drives that provide cryptographic erase such as ISE and SED SAS and SATA drives, and PCIe SSDs.
- Overwrite erase drives— Includes all drives that do not support cryptographic erase.

(i) NOTE: Before erasing vFlash, you must first detach all partitions using iDRAC interfaces before executing the operation.

NOTE: System erase only applies to drives within the server. iDRAC is not able to erase drives in an external enclosure such as a JBOD.

The RACADM SystemErase sub-command includes options for the following categories:

- The **SecureErasePD** option cryptographically erases all the secure erase drives.
- The OverwritePD option overwrites data on all drives.
- () NOTE: Cryptographic Erase of BOSS physical disk can be done by SystemErase method and it is supported from LC-UI, Wsman, and Racadm

Before performing SystemErase, use the following command to check the erase capability of all physical disks for a server:

racadm storage get pdisks -o -p SystemEraseCapability

() NOTE: If SEKM is enabled on the server, then disable SEKM using the racadm sekm disable command before using this command. This can avoid any storage devices being locked out which are secured by iDRAC, if SEKM settings are erased from iDRAC by executing this command.

To erase ISE and SED drives, use this command:

racadm systemerase -secureerasepd

To erase overwrite erase drives, use the following command:

racadm systemerase -overwritepd

(i) NOTE: RACADM SystemErase removes all the virtual disks from the physical disks that are erased by the above commands.

(i) NOTE: RACADM SystemErase causes the server to restart in order to perform the erase operations.

NOTE: Individual PCIe SSD or SED devices can be erased using the iDRAC GUI or RACADM. For more information, see the *Erasing PCIe SSD device data* and the *Erasing SED device data* section.

For information on the System Erase function within the Lifecycle Controller GUI, see the *Lifecycle Controller User's Guide* available at https://www.dell.com/idracmanuals .

Erasing SED/ISE device data

NOTE: This operation is not supported when supported device is a part of a Virtual Disk. The target supported device must be removed from the virtual disk prior to performing device erase.

Cryptographic Erase permanently erases all data present on the disk. Performing a Cryptographic Erase on an SED/ISE overwrites all blocks and results in permanent loss of all data on the supported devices. During Cryptographic Erase, the host is

unable to access the supported device. SED/ISE device erase can be performed either in real time or be applied after a system reboot.

Fls.<u>1282</u> Mov. 33

If the system reboot or experiences a power loss during cryptographic erase, the operation is canceled. You must reboot the system and restart the process.

Before erasing SED/ISE device data, ensure that:

- Lifecycle Controller is enabled.
- You have Server Control and Login privileges.
- Selected supported drive is not part of a virtual disk.

() NOTE:

- Erasing SED/ISE can be performed either as a real time or as a staged operation.
- After the drive is erased, it may still be displayed as active within the OS due to data caching. If this occurs, reboot the OS and the erased drive will no longer be displayed or report any data.
- Cryptographic erase operation is not supported for hot-plugged NVMe disks. Reboot the server before starting the operation. If the operation continues to fail, ensure that CSIOR is enabled and that the NVMe disks qualified by Dell Technologies.

Erasing SED/ISE device data using web interface

To erase the data on the supported device:

- 1. In the iDRAC Web interface, go to **Storage** > **Overview** > **Physical Disks**.
 - The **Physical Disk** page is displayed.
- 2. From the **Controller** drop-down menu, select the controller to view the associated devices.
- 3. From the drop-down menus, select Cryptographic Erase for one or more SED/ISEs.

If you have selected **Cryptographic Erase** and you want to view the other options in the drop-down menu, then select **Action** and then click the drop-down menu to view the other options.

- 4. From the Apply Operation Mode drop-down menu, select one of the following options:
 - Apply Now Select this option to apply the actions immediately with no system reboot required.
 - At Next Reboot Select this option to apply the actions during the next system reboot.
 - At Scheduled Time Select this option to apply the actions at a scheduled day and time:
 - **Start Time** and **End Time** Click the calendar icons and select the days. From the drop-down menus, select the time. The action is applied between the start time and end time.
 - From the drop-down menu, select the type of reboot:
 - No Reboot (Manually Reboot System)
 - Graceful Shutdown
 - Force Shutdown
 - Power Cycle System (cold boot)
- 5. Click Apply.

If the job is not created, a message indicating that the job creation was not successful is displayed. Also, the message ID and the recommended response action is displayed.

If the job is created successfully, a message indicating that the job ID is created for the selected controller is displayed. Click **Job Queue** to view the progress of the job in the Job Queue page.

If pending operation is not created, an error message is displayed. If pending operation is successful and job creation is not successful, then an error message is displayed.

Erasing SED device data using RACADM

To securely erase an SED device:

```
racadm storage cryptographicerase:<SED FQDD>
```

FIS.1283 Mov. 33

To create the target job after executing the cryptographicerase command:

racadm jobqueue create <SED FQDD> -s TIME NOW -realtime

To create the target staged job after executing the cryptographicerase command:

racadm jobqueue create <SED FQDD> -s TIME NOW -e <start time>

To query the job ID returned:

racadm jobqueue view -i <job ID>

For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Rebuild Physical Disk

Rebuild Physical Disk is the ability to reconstruct the contents of a failed disk. This is true only when auto rebuild option is set to false. If there is a redundant virtual disk, the rebuild operation can reconstruct the contents of a failed physical disk. A rebuild can take place during normal operation, but it degrades performance.

Cancel Rebuild can be used to cancel a rebuild that is in progress. If you cancel a rebuild, the virtual disk remains in a degraded state. The failure of an additional physical disk can cause the virtual disk to fail and may result in data loss. It is recommended to perform a rebuild on the failed physical disk at the earliest.

In case, you cancel the rebuild of a physical disk that is assigned as a hot spare, reinitiate the rebuild on the same physical disk in order to restore the data. Canceling the rebuild of a physical disk and then assigning another physical disk as a hot spare does not cause the newly assigned hot spare to rebuild the data.

Managing virtual disks

You can perform the following operations for the virtual disks:

- Create
- Delete
- Edit policies
- Initialize
- Check consistency
- Cancel check consistency
- Encrypt virtual disks
- Assign or unassign dedicated hot spares
- Blink and unblink virtual disk
- Cancel background initialization
- Online capacity expansion
- RAID level migration

NOTE: You can manage and monitor 240 virtual disks using iDRAC interfaces. To create VDs, use either Device Setup (F2), PERCCLI command line tool, or Dell OpenManage Server Administrator (OMSA).

(i) NOTE: PERC 10 count is less since it does not support daisy chain arrangements.

Creating virtual disks

To implement RAID functions, you must create a virtual disk. A virtual disk refers to storage created by a RAID controller from one or more physical disks. Although a virtual disk may be created from several physical disks, it is seen by the operating system as a single disk.

Before creating a virtual disk, you should be familiar with the information in Considerations Before Creating Virtual Disks.

You can create a Virtual Disk using the Physical Disks attached to the PERC controller. To create a Virtual Disk, you must have the Server Control user privilege. You can create a maximum of 64 virtual drives and a maximum of 16 virtual drives in the same drive group.



You cannot create a virtual disk if:

- Physical disk drives are not available for virtual disk creation. Install additional physical disk drives.
- Maximum number of virtual disks that can be created on the controller has been reached. You must delete at least one virtual disk and then create a new virtual disk.
- Maximum number of virtual disks supported by a drive group has been reached. You must delete one virtual disk from the selected group and then create a new virtual disk.
- A job is currently running or scheduled on the selected controller. You must wait for this job to complete or you can delete the job before attempting a new operation. You can view and manage the status of the scheduled job in the Job Queue page.
- Physical disk is in non-RAID mode. You must convert to RAID mode using iDRAC interfaces such as iDRAC web interface, RACADM, Redfish, WSMan, or <CTRL+R>.

NOTE: If you create a virtual disk in Add to Pending Operation mode and a job is not created, and then if you delete the Virtual disk, then the create pending operation for the virtual disk is cleared.

(i) NOTE: RAID 6 and 60 are not supported in PERC H330.

() NOTE: BOSS controller allows you to create virtual disk only of size equal to the full size of the M.2 physical storage media. Ensure that you set the virtual disk size to zero when using the Server Configuration Profile to create a BOSS virtual disk. For other interfaces such as RACADM, WSMan, and Redfish, the virtual disk size should not be specified.

Considerations before creating virtual disks

Before creating virtual disks, consider the following:

- Virtual disk names not stored on controller—The names of the virtual disks that you create are not stored on the controller. This means that if you reboot using a different operating system, the new operating system may rename the virtual disk using its own naming conventions.
- Disk grouping is a logical grouping of disks attached to a RAID controller on which one or more virtual disks are created, such that all virtual disks in the disk group use all of the physical disks in the disk group. The current implementation supports the blocking of mixed disk groups during the creation of logical devices.
- Physical disks are bound to disk groups. Therefore, there is no RAID level mixing on one disk group.
- There are limitations on the number of physical disks that can be included in the virtual disk. These limitations depend on the controller. When creating a virtual disk, controllers support a certain number of stripes and spans (methods for combining the storage on physical disks). Because the number of total stripes and spans is limited, the number of physical disks that can be used is also limited. The limitations on stripes and spans affect the RAID levels as follows:
 - Maximum number of spans affects RAID 10, RAID 50, and RAID 60.
 - Maximum number of stripes affects RAID 0, RAID 5, RAID 50, RAID 6, and RAID 60.
 - \circ $\,$ Number of physical disks in a mirror is always 2. This affects RAID 1 and RAID 10.
 - (i) NOTE:
 - RAID 1 is only supported for BOSS controllers.
 - SWRAID controller only supports RAID 0, 1, 5 and 10.
- Cannot create virtual disks on PCIe SSDs. But PERC 11 and later controllers support creating virtual disks using PCIe SSDs.

Creating virtual disks using web interface

To create virtual disk:

- 1. In the iDRAC Web interface, go to **Storage** > **Overview** > **Virtual DisksAdvanced Filter**.
- 2. In the Virtual Disk section, do the following:
 - a. From the Controller drop-down menu, select the controller for which you want to create the virtual disk.
 - b. From the Layout drop-down menu, select the RAID level for the Virtual Disk.

Only those RAID levels supported by the controller appear in the drop-down menu and it is based on the RAID levels are available based on the total number of physical disks available.

- c. Select the Media Type, Stripe Size, Read Policy, Write Policy, Disk Cache Policy, .
 Only those values supported by the controller appear in the drop-down menus for these properties.
- d. In the Capacity field, enter the size of the virtual disk.
- The maximum size is displayed and then updated as disks are selected.
- e. The **Span Count** field is displayed based on the selected physical disks (step 3). You cannot set this value. It is automatically calculated after selecting disks for multi-raid level. **Span Count** field is applicable to RAID 10, RAID 50, and



RAID 60. If you have selected RAID 10 and if the controller supports uneven RAID 10, then the span count value is not displayed. The controller automatically sets the appropriate value. For RAID 50 and RAID 60, this field is not displayed when minimum number of disks are used to create RAID. It can be changed if more disks are used.

- **3.** In the **Select Physical Disks** section, select the number of physical disks. For more information about the fields, see the *iDRAC Online Help*
- 4. From the Apply Operation Mode drop-down menu, select when you want to apply the settings.
- 5. Click Create Virtual Disk.

Based on the selected **Apply Operation Mode**, the settings are applied.

NOTE: You can use alphanumeric characters, spaces, dashes, and underscores in the disk name.

Any other special characters that you enter are removed and replaced by space while creating the virtual disk.

Creating virtual disks using RACADM

Use the racadm storage createvd command.

For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

(i) NOTE: Disk slicing or configuring partial VDs is not supported using RACADM on the drives managed by S140 controller.

Editing virtual disk cache policies

You can change the read, write, or disk cache policy of a virtual disk.

NOTE: Some of the controllers do not support all read or write policies. Therefore, when a policy is applied, an error message is displayed.

The read policies indicate whether the controller must read sequential sectors of the virtual disk searching for data:

- Adaptive Read Ahead The controller initiates read ahead only if the two most recent reads requests accessed sequential sectors of the disk. If subsequent read requests access random sectors of the disk, the controller reverts to no read ahead policy. The controller continues to evaluate whether read requests are accessing sequential sectors of the disk, and initiates read ahead if necessary.
- **Read Ahead** The controller reads sequential sectors of the virtual disk when seeking data. Read ahead policy may improve system performance if the data is written to the sequential sectors of the virtual disk.
- No Read Ahead Selecting no read ahead policy indicates that the controller should not use read ahead policy.

The write policies specify if the controller sends a write-request completion signal when the data is in the cache or after it has been written to the disk.

- Write Through The controller sends a write-request completion signal only after the data is written to the disk. Write-through caching provides better data security than write-back caching, since the system assumes that the data is available only after it has been safely written to the disk.
- Write Back The controller sends a write-request completion signal as soon as the data is in the controller cache but has not yet been written to disk. Write back caching may provide improved performance since subsequent read requests can retrieve data quickly from the cache then from the disk. However, data loss may occur in the event of a system failure which prevents that data from being written on a disk. Other applications may also experience problems when actions assume that the data is available on the disk.
- Force Write Back The write cache is enabled regardless of whether the controller has a battery. If the controller does not have a battery and force write-back caching is used, data loss may occur in the event of a power failure.

The Disk Cache policy applies to readings on a specific virtual disk. These settings do not affect the read-ahead policy.

() NOTE:

- Controller non-volatile cache and battery backup of controller cache affects the read-policy or the write policy that a controller can support. All PERCs do not have battery and cache.
- Read ahead and write back requires cache. Therefore, if the controller does not have cache, it does not allow you to set the policy value.

Similarly, if the PERC has cache but not battery and the policy is set that requires accessing cache, then data loss may occur if base of power off. So few PERCs may not allow that policy.

Therefore, depending upon the PERC, the policy value is set.



Deleting virtual disks

Deleting a virtual disk destroys all information including file systems and volumes residing on the virtual disk and removes the virtual disk from the controller's configuration. When deleting virtual disks, all assigned global hot spares may be automatically unassigned when the last virtual disk associated with the controller is deleted. When deleting the last virtual disk of a disk group, all assigned dedicated hot spares automatically become global hot spares.

If you delete all the VDs for a global hotspare, then the global hotspare gets automatically deleted.

You must have the Login and Server Control privilege to perform delete virtual disks.

When this operation is allowed, you can delete a boot virtual drive. It is done from sideband and the independent of the operating system. Hence, a warning message appears before you delete the virtual drive.

If you delete a virtual disk and immediately create a new virtual disk with all the same characteristics as the one that was deleted, the controller recognizes the data as if the first virtual disk were never deleted. In this situation, if you do not want the old data after recreating a new virtual disk, re-initialize the virtual disk.

Checking virtual disk consistency

This operation verifies the accuracy of the redundant (parity) information. This task only applies to redundant virtual disks. When necessary, the check consistency task rebuilds the redundant data. If the virtual drive has a degraded status, running a check consistency may be able to return the virtual drive to ready status. You can perform a consistency check using the web interface or RACADM.

You can also cancel the check consistency operation. The cancel check consistency is a real-time operation.

You must have Login and Server Control privilege to check consistency of virtual disks.

(i) NOTE: Consistency check is not supported when the drives are set up in RAIDO mode.

NOTE: If you perform Cancel Consistency operation when there is no consistency check operation is in progress, then the pending operation in GUI is shown as Cancel BGI instead of Cancel Consistency check.

Initializing virtual disks

Initializing virtual disks erases the all the data on the disk but does not change the virtual disk configuration. You must initialize a virtual disk that is configured before it is used.

(i) NOTE: Do not initialize virtual disks when attempting to recreate an existing configuration.

You can perform a fast initialization, a full Initialization, or cancel the initialization operation.

NOTE: The cancel initialization is a real-time operation. You can cancel the initialization using only the iDRAC Web interface and not RACADM.

Fast initialization

The fast initialize operation initializes all physical disks included in the virtual disk. It updates the metadata on the physical disks so that all disk space is available for future write operations. The initialize task can be completed quickly because the existing information on the physical disks is not erased, although future write operations overwrite any information that remains on the physical disks.

Fast initialization only deletes the boot sector and stripe information. Perform a fast initialize only if you are constrained for time or the hard drives are new or unused. Fast Initialization takes less time to complete (usually 30-60 seconds).

CAUTION: Performing a fast initialize causes existing data to be inaccessible.

The fast initialize task does not write zeroes to the disk blocks on the physical disks. It is because the Fast Initialize task does not perform a write operation, it causes less degradation to the disk.

244 Managing storage devices



A fast initialization on a virtual disk overwrites the first and last 8 MB of the virtual disk, clearing any boot records or partition information. The operation takes only 2-3 seconds to complete and is recommended when you are recreating virtual disks.

A background initialization starts five minutes after the Fast Initialization is completed.

Full or slow initialization

The full initialization (also called slow initialize) operation initializes all physical disks included in the virtual disk. It updates the metadata on the physical disks and erases all existing data and file systems. You can perform a full initialization after creating the virtual disk. In comparison with the fast initialize operation, you may want to use the full initialize if you have trouble with a physical disk or suspect that it has bad disk blocks. The full initialize operation remaps bad blocks and writes zeroes to all disk blocks.

If full initialization of a virtual disk is performed, background initialization is not required. During full initialization, the host is not able to access the virtual disk. If the system reboots during a full initialization, the operation terminates and a background initialization process starts on the virtual disk.

It is always recommended to do a full initialization on drives that previously contained data. Full initialization can take up to 1-2 minutes per GB. The speed of initialization depends on the controller model, speed of hard drives, and the firmware version.

The full initialize task initializes one physical disk at a time.

(i) NOTE: Full initialize is supported only in real-time. Only few controllers support full initialization.

Encrypting virtual disks

When encryption is disabled on a controller (that is, the security key is deleted), manually enable encryption for virtual disks created using SED drives. If the virtual disk is created after encryption is enabled on a controller, the virtual disk is automatically encrypted. It is automatically configured as an encrypted virtual disk unless the enabled encryption option is disabled during the virtual disk creation.

You must have Login and Server Control privilege to manage the encryption keys.

NOTE: Though encryption is enabled in the controllers, user needs to manually enable encryption on the VD if VD is created from iDRAC. Only if the VD is created from OMSA, it would be automatically encrypted.

Assigning or unassigning dedicated hot spares

A dedicated hot spare is an unused backup disk that is assigned to a virtual disk. When a physical disk in the virtual disk fails, the hot spare is activated to replace the failed physical disk without interrupting the system or requiring your intervention.

You must have Login and Server Control privilege to run this operation.

You can assign only 4K drives as hot spare to 4K virtual disks.

If you have assigned a physical disk as a dedicated hot spare in Add to Pending Operation mode, the pending operation is created but a job is not created. Then, if you try to unassign the dedicated hot spare, the assign dedicated hot spare pending operation is cleared.

If you have unassigned a physical disk as a dedicated hot spare in Add to Pending Operation mode, the pending operation is created but a job is not created. Then, if you try to assign the dedicated hot spare, the unassign dedicated hot spare pending operation is cleared.

NOTE: While the log export operation is in progress, you cannot view information about dedicated hot spares on the

Manage Virtual Disks page. After the log export operation is complete, reload or refresh the **Manage Virtual Disks** page to view the information.

Rename VD

To change the name of a Virtual Disk, the user must have System Control privilege. The virtual disk name can contain only alphanumeric characters, spaces, dashes and underscores. The maximum length of the name depends on the individual controller. In most cases, the maximum length is 15 characters. The name cannot start with a space, end with a space, or be left blank. Every time a virtual disk is renamed, an LC Log gets created.



Edit Disk capacity

Online Capacity Expansion (OCE) allows you to increase the storage capacity of selected RAID levels while the system remains online. The controller redistributes the data on the array(called Reconfiguration), placing new space available at the end of each RAID array.

Online Capacity Expansion (OCE) can be achieved in two ways:

- If free space is available on the smallest physical drive on the virtual disks group after starting LBA of Virtual disks, the virtual disk's capacity can be expanded within that free space. This option allows you to enter the new increased virtual disk size. If disk group in a virtual disk has space available only before starting LBA, then Edit Disk Capacity in same disk group is not permitted even though there is Available Space on a physical drive.
- A virtual disk's capacity can also be expanded by adding additional compatible physical disks to the existing virtual disk group. This option does not allow you to enter the new increased virtual disk size. New increased virtual disk size is calculated and displayed to the user based on the used disk space of existing physical disk group on a particular virtual disk, existing raid level of the virtual disk and the number of new drives added to the virtual disk.

Capacity Expansion allows user to specify the final VD size. Internally final VD size is conveyed to PERC in percentage (this percentage is the space user would like to use from empty space left in the array for the local disk to expand). Because of this percentage logic final VD size after reconfiguration completes may be different from what user provided for scenario where user is not giving maximum VD size possible as the final VD size (percentage turns out to be less than 100%). User does not see difference in this entered VD size and final VD size after reconfiguration, if maximum possible VD size is entered by user.

Raid Level Migration

RAID Level Migration (RLM) refers to changing a virtual disk's RAID level. iDRAC9 provides an option to increase the VD size using RLM. In a way, RLM allows migrating the RAID level of a virtual disk which in turn may increase the size of virtual disk.

RAID level migration is the process of converting a VD with one RAID Level to another. When you migrate a VD to a different Raid Level, the user data on it is redistributed to the format of the new configuration.

This configuration is supported by both staged and realtime.

The below table describes possible reconfigurable VD layouts while reconfiguring (RLM) a VD with addition of disks and without addition of disks.

Table 54. Possible VD Layout

Source VD Layout	Possible target VD Layout with Disk Add	Possible target VD Layout Without disk addition
R0 (single disk)	R1	NA
RO	R5/R6	NA
R1	R0/R5/R6	RO
R5	R0/R6	RO
R6	R0/R5	R0/R5

Permitted operations when OCE or RLM is going on

The following operations are allowed when OCE/RLM is going on:

Table 55. Permitted operations

From Controller End behind which a VD is going through OCE/RLM	From VD End (which is going through OCE/RLM)	From any other Ready State Physical Disk on the same controller	From any other VD (which is not going through OCE/ RLM) End on the same controller
Reset Configuration	Delete	Blink	Delete
Export Log	Blink	Unblink	Blink



Table 55. Permitted operations (continued)

From Controller End behind which a VD is going through OCE/RLM	From VD End (which is going through OCE/RLM)	From any other Ready State Physical Disk on the same controller	From any other VD (which is not going through OCE/ RLM) End on the same controller
Set Patrol Read Mode	Unblink	Assign Global Hot Spare	Unblink
Start Patrol Read		Convert to non-RAID Disks	Rename
Change Controller Properties			Change Policy
Manage Physical Disk Power			Slow Initialize
Convert to RAID Capable Disks			Fast Initialize
Convert to Non-RAID Disks			Replace Member Disk
Change Controller Mode			

OCE and RLM Restrictions or Limitations

Following are the common limitations for OCE and RLM:

- OCE/RLM is restricted to the scenario where the disk group contains only one VD.
- OCE is not supported on RAID50 and RAID60. RLM is not supported on RAID10, RAID50 and RAID60.
- If the controller already contains the maximum number of virtual disks, you cannot perform a RAID level migration or capacity expansion on any virtual disk.
- The controller changes the write cache policy of all virtual disks undergoing a RLM/OCE to Write-Through until RLM/OCE is complete.
- Reconfiguring Virtual Disks typically impacts disk performance until the reconfiguration operation is complete.
- The total number of physical disks in a disk group cannot exceed 32.
- If any background operation (like BGI/rebuild/copyback/patrol read) is already running on the corresponding VD/PD then Reconfiguration (OCE/RLM) is not allowed at that time.
- Any kind of disk migration when Reconfiguration (OCE/RLM) is on progress on drives associated with VD causes reconfiguration to fail.
- Any new drive added for OCE/RLM becomes part of the VD after reconstruction completes. But State for those new drive changes to Online just after reconstruction starts.

Cancel Initialization

This feature is the ability to cancel the background initialization on a virtual disk. On PERC controllers, the background initialization of redundant virtual disk starts automatically after a virtual disk is created. The background initialization of redundant virtual disk prepares the virtual disk for parity information and improves write performance. However, some processes such as creating a virtual disk cannot be run while the background initialization is in progress. Cancel Initialization provides the ability to cancel the background initialization manually. Once cancelled, the background initialization automatically restarts within 0 to 5 minutes.

(i) NOTE: Background initialization is not applicable for RAID 0 virtual disks.

Managing virtual disks using web interface

- 1. In the iDRAC web interface, go to Configuration > Storage Configuration > Virtual Disk Configuration.
- 2. From the Virtual Disks, select the controller for which you want to manage the virtual disks.
- 3. From Action drop-down menu, select an action.
 - When you select an action, an additional Action window displayed. Select / enter the desired value.
 - Rename
 - Delete



- Edit Cache Policy You can change the cache policy for the following options:
 - **Read Policy** Following values are available for selection:
 - Adaptive Read Ahead Indicates that for the given volume, the control uses the Read-Ahead cache policy if the two most recent disks accesses occurred in sequential sectors. If the read requests are random, the controller returns to No Read Ahead mode.
 - No Read Ahead Indicates that for the given volume, no read ahead policy is used.
 - Read Ahead Indicates that for the given volume, the controller reads sequentially ahead of the requested data and stores the additional data in cache memory, anticipating a data requirement. This speeds up sequential data reads, but there is less improvement when accessing random data.
 - Write Policy Change the write cache policy to one of the following options:
 - Write Through Indicates that for the given volume, the controller sends a data transfer completion signal to the host system when the disk subsystem has received all the data in a transaction.
 - Write Back Indicates that for the given volume, the controller sends a data transfer completion signal to the host system when the controller cache has received all the data in a transaction. The controller then writes the cached data to the storage device in the background.
 - Force Write Back When using force write-back caching, the write cache is enabled regardless of whether the controller has a battery. If the controller does not have a battery and force write-back caching is used, data loss may occur in the event of a power failure.
 - **Disk Cache Policy** Change the disk cache policy to one of the following options:
 - **Default** Indicates that the disk is using its default write cache mode. For SATA disks, this is enabled and for SAS disks this is disabled.
 - **Enabled** Indicates that the disk's write cache is enabled. This increases performance and the probability of data loss if there is power loss.
 - Disabled Indicates that the disk's write cache is disabled. This decreases performance and the probability of data loss.
- Edit Disk Capacity You can add the physical disks to the selected virtual disk in this window. This window also shows the current capacity and new capacity of the virtual disk after adding the physical disks.
- **RAID Level Migration** Displays the Disk Name, Current RAID Level, and size of the virtual disk. Allows you to select a New RAID Level. User may have to add additional drives to existing Virtual disks to migrate to new raid level. This feature is not applicable on RAID 10, 50 and 60.
- Initialize: Fast Updates the metadata on the physical disks so that all the disk space is available for future write operations. The initialize option can be completed quickly because existing information on the physical disks is not erased, although future write operations overwrites any information that remains on the physical disks.
- Initialize: Full All existing data and file systems are erased.

(i) NOTE: The Initialize: Full option is not applicable for PERC H330 controllers.

• Check Consistency — To check the consistency of a virtual disk, select Check Consistency from the corresponding drop-down menu.

(i) NOTE: Consistency check is not supported on drives set up in RAID0 mode.

For more information about these options, see the *iDRAC Online Help*.

 Click Apply Now to apply the changes immediately, At Next Reboot to apply the changes after next reboot, At Scheduled Time to apply the changes at a particular time, and Discard All Pending to discard the changes. Based on the selected operation mode, the settings are applied.

Managing virtual disks using RACADM

Use the following commands to manage virtual disks:

• To delete virtual disk:

racadm storage deletevd:<VD FQDD>

To initialize virtual disk:

racadm storage init:<VD FQDD> -speed {fast|full}

To check consistency of virtual disks (not supported on RAID0):

```
racadm storage ccheck: <vdisk fqdd>
```

To cancel the consistency check:



racadm storage cancelcheck: <vdisks fqdd>

• To encrypt virtual disks:

racadm storage encryptvd:<VD FQDD>

• To assign or unassign dedicated hot spares:

```
racadm storage hotspare:<Physical Disk FQDD> -assign <option> -type dhs -vdkey: <FQDD
of VD>
```

<option>=yes

Assign hot spare

<Option>=NO

Unassign hot spare

RAID Configuration Features

Following table lists some of the RAID configuration features which are available in RACADM and WSMan:

CAUTION: Forcing a physical disk to go online or offline may result in data loss.

Table 56. RAID Configuration Features

Feature	RACADM Command	Description	
Force Online	racadm storage forceonline: <pd fqdd=""></pd>	A power failure, corrupted data, or some other reason may lead to a physical disk going offline. You can use this feature to force a physical disk back into an online state when all other options have been exhausted. Once the command is run, the controller places the drive back into online state and restore its membership within the virtual disk. This happens only if the controller can read from the drive and can write into its metadata.	
() NOTE: Data recovery is only possible if a limited portion of the disk is damaged. Force Online feature cannot fix an already failed disk.			
Force Offline	racadm storage forceoffline: <pd fqdd=""></pd>	This feature removes a drive from a virtual disk configuration so that it goes offline, resulting in a degraded VD configuration. It is helpful if a drive is likely to fail in near future or is reporting a SMART failure but is still online. It can be also used if you would like to utilize a drive which is part of an existing RAID configuration.	
Replace Physical Disk	racadm storage replacephysicaldisk: <source PD FQDD > -dstpd <destination fqdd="" pd=""></destination></source 	Allows you to copy data from a physical disk which is a member of a VD, to another physical disk. The source disk should be in online state, while the destination disk should be in ready state and of a similar size and type to replace the source.	
Virtual Disk as boot device	racadm storage setbootvd: <controller< td=""><td>A virtual disk can be configured as a boot device using this feature. This enables fault tolerance when a VD with</td></controller<>	A virtual disk can be configured as a boot device using this feature. This enables fault tolerance when a VD with	



Table 56. RAID Configuration Features (continued)

Feature	RACADM Command	Description
	FQDD> -vd <virtualdisk FQDD></virtualdisk 	redundancy is selected as the boot device, and also has the operating system installed on it.
Unlock Foreign Configuration	racadm storage unlock: <controller fqdd=""> -key <key id=""> -passwd <passphrase></passphrase></key></controller>	This feature is used to authenticate locked drives which have a different source controller encryption than the destination. Once unlocked, the drive can be successfully migrated from one controller to another.

Managing controllers

You can perform the following for controllers:

- Configure controller properties
- Import or auto import foreign configuration
- Clear foreign configuration
- Reset controller configuration
- Create, change, or delete security keys
- Discard preserved cache

Configuring controller properties

You can configure the following properties for the controller:

- Patrol read mode (auto or manual)
- Start or stop patrol read if patrol read mode is manual
- Patrol read unconfigured areas
- Check consistency mode
- Copyback mode
- Load balance mode
- Check consistency rate
- Rebuild rate
- BGI rate
- Reconstruct rate
- Enhanced auto import foreign configuration
- Create or change security keys
- Encryption mode (Local Key Management and Secure Enterprise key Manager)

You must have Login and Server Control privilege to configure the controller properties.

Patrol read mode considerations

Patrol read identifies disk errors to avoid disk failures, data loss, or corruption. It runs automatically once a week on SAS and SATA HDDs.

The Patrol Read does not run on a physical disk in the following circumstances:

- The physical disk is an SSD.
- The physical disk is not included in a virtual disk or assigned as a hot spare.



- The physical disk is included in a virtual disk that is undergoing one of the following:
 - A rebuild
 - A reconfiguration or reconstruction
 - A background initialization
 - A check consistency

In addition, the Patrol Read operation suspends during heavy I/O activity and resumes when the I/O is complete.

NOTE: For more information on how often the Patrol Read operation runs when in auto mode, see the respective controller documentation.

() NOTE: Patrol read mode operations such as **Start** and **Stop** are not supported if there are no virtual disks available in the controller. Though you can invoke the operations successfully using the iDRAC interfaces, the operations fail when the associated job is started.

Load balance

The Load Balance property provides the ability to automatically use both controller ports or connectors connected to the same enclosure to route I/O requests. This property is available only on SAS controllers.

Bgi rate

(i) NOTE: Both H330 and H345 require the driver to be loaded for the background initialization operations to run.

On PERC controllers, background initialization of a redundant virtual disk begins automatically within 0 to 5 minutes after the virtual disk is created. The background initialization of a redundant virtual disk prepares the virtual disk to maintain redundant data and improves write performance. For example, after the background initialization of a RAID 5 virtual disk completes, the parity information has been initialized. After the background initialization of a RAID 1 virtual disk completes, the physical disks are mirrored.

The background initialization process helps the controller identify and correct problems that may occur with the redundant data later. In this regard, the background initialization process is similar to a check consistency. The background initialization should be allowed to run to completion. If cancelled, the background initialization automatically restarts within 0 to 5 minutes. Some processes such as read and write operations are possible while the background initialization is running. Other processes, such as creating a virtual disk, cannot be run concurrently with a background initialization. These processes cause the background initialization to cancel.

The background initialization rate, configurable between 0% and 100%, represents the percentage of the system resources dedicated to running the background initialization task. At 0%, the background initialization has the lowest priority for the controller, takes the most time to complete, and is the setting with the least impact to system performance. A background initialization rate of 0% does not mean that the background initialization is stopped or paused. At 100%, the background initialization is the highest priority for the controller. The background initialization time is minimized and is the setting with the most impact to system performance.

Check consistency

The Check Consistency task verifies the accuracy of the redundant (parity) information. This task only applies to redundant virtual disks. When necessary, the Check Consistency task rebuilds the redundant data. If the virtual disk is in a Failed Redundancy state, running a check consistency may be able to return the virtual disk to a Ready state.

The check consistency rate, configurable between 0% and 100%, represents the percentage of the system resources dedicated to running the check consistency task. At 0%, the check consistency has the lowest priority for the controller, takes the most time to complete, and is the setting with the least impact to system performance. A check consistency rate of 0% does not mean that the check consistency is stopped or paused. At 100%, the check consistency is the highest priority for the controller. The check consistency time is minimized and is the setting with the most impact to system performance.

Create or change security keys

When configuring the controller properties, you can create or change the security keys. The controller uses the encryption key to lock or unlock access to SED. You can create only one encryption key for each encryption-capable controller. The security key is managed using following features:



- Local Key Management (LKM) System LKM is used to generate the key ID and the password or key required to secure the virtual disk. If you are using LKM, you must create the encryption key by providing the Security Key Identifier and the Passphrase.
- 2. Secure Enterprise Key Manager (SEKM) This feature is used to generate the key using the Key Management Server (KMS). If you are using SEKM, you must configure iDRAC with KMS information as well as SSL related configuration.

() NOTE:

- This task is not supported on PERC hardware controllers running in eHBA mode.
- If you create the security key in 'Add to Pending Operation' mode and a job is not created, and then if you delete the security key, the create security key pending operation is cleared.

() NOTE:

- For enabling SEKM, ensure that the supported PERC firmware is installed.
- You cannot downgrade the PERC firmware to the previous version if SEKM is enabled. Downgrading of other PERC controller firmware in the same system which is not in SEKM mode may also fail. To downgrade the firmware for the PERC controllers that are not in SEKM mode, you can use OS DUP update method, or disable SEKM on the controllers and then retry the downgrade from iDRAC.

NOTE: When importing a hot plugged locked volume from one server to another, you will see CTL entries for Controller attributes being applied in the LC Log.

Configuring controller properties using web interface

- 1. In the iDRAC web interface, go to **Storage** > **Overview** > **Controllers**. The **Setup Controllers** page is displayed.
- 2. In the Controller section, select the controller that you want to configure.
- **3.** Specify the required information for the various properties.

The **Current Value** column displays the existing values for each property. You can modify this value by selecting the option from the **Action** drop-down menu for each property.

For information about the fields, see the *iDRAC Online Help*.

- 4. From the Apply Operation Mode, select when you want to apply the settings.
- 5. Click Apply.

Based on the selected operation mode, the settings are applied.

Configuring controller properties using RACADM

• To set Patrol Read Mode:

racadm set storage.controller.<index>.PatrolReadMode {Automatic | Manual | Disabled}

• If Patrol read mode is set to manual, use the following commands to start and stop Patrol read Mode:

racadm storage patrolread:<Controller FQDD> -state {start|stop}

- () NOTE: Patrol read mode operations such as Start and Stop are not supported if there are no virtual disks available in the controller. Though you can invoke the operations successfully using the iDRAC interfaces, the operations will fail when the associated job is started.
- To specify the Check Consistency Mode, use **Storage.Controller.CheckConsistencyMode** object.
- To enable or disable the Copyback Mode, use Storage.Controller.CopybackMode object.
- To enable or disable the Load Balance Mode, use Storage.Controller.PossibleloadBalancedMode object.
- To specify the percentage of the system's resources dedicated to perform a check consistency on a redundant virtual disk, use **Storage.Controller.CheckConsistencyRate** object.
- To specify the percentage of the controller's resources dedicated to rebuild a failed disk, use Storage.Controller.RebuildRate object



- To specify the percentage of the controller's resources dedicated to perform the background initialization (BGI) of a virtual disk after it is created, use**Storage.Controller.BackgroundInitializationRate** object
- To specify the percentage of the controller's resources dedicated to reconstruct a disk group after adding a physical disk or changing the RAID level of a virtual disk residing on the disk group, use **Storage.Controller.ReconstructRate** object
- To enable or disable the enhanced auto import of foreign configuration for the controller, use **Storage.Controller.EnhancedAutoImportForeignConfig** object
- To create, modify, or delete security key to encrypt virtual drives:

```
racadm storage createsecuritykey:<Controller FQDD> -key <Key id> -passwd <passphrase>
racadm storage modifysecuritykey:<Controller FQDD> -key <key id> -oldpasswd <old
passphrase> -newpasswd <new passphrase>
racadm storage deletesecuritykey:<Controller FQDD>
```

Importing or auto importing foreign configuration

A foreign configuration is data residing on physical disks that have been moved from one controller to another. Virtual disks residing on physical disks that have been moved are considered to be a foreign configuration.

You can import foreign configurations so that virtual disks are not lost after moving Physical Disks. A foreign configuration can be imported only if it contains a virtual disk that is in either Ready or Degraded state or a hotspare that is dedicated to a virtual disk which can be imported or is already present.

All of the virtual disk data must be present, but if the virtual disk is using a redundant RAID level, the additional redundant data is not required.

For example, if the foreign configuration contains only one side of a mirror in a RAID 1 virtual disk, then the virtual disk is in a Degraded state and can be imported. If the foreign configuration contains only one physical disk that was originally configured as a RAID 5 using three physical disks, then the RAID 5 virtual disk is in a Failed state and cannot be imported.

In addition to virtual disks, a foreign configuration may consist of a physical disk that was assigned as a hot spare on one controller and then moved to another controller. The Import Foreign Configuration task imports the new physical disk as a hot spare. If the physical disk was set as a dedicated hot spare on the previous controller, but the virtual disk to which the hot spare was assigned is no longer present in the foreign configuration, then the physical disk is imported as a global hot spare.

If any foreign configurations locked using Local Key manager (LKM) are Detected, then import foreign configuration operation is not possible in iDRAC in this release. You must unlock the drives through CTRL-R and then continue to import foreign configuration from iDRAC.

The Import Foreign Configuration task is only displayed when the controller has detected a foreign configuration. You can also identify whether a physical disk contains a foreign configuration (virtual disk or hot spare) by checking the physical disk state. If the physical disk state is Foreign, then the physical disk contains all or some portion of a virtual disk or has a hot spare assignment.

() NOTE: The task of importing foreign configuration imports all virtual disks residing on physical disks that have been added to the controller. If more than one foreign virtual disk is present, all the configurations are imported.

PERC9 controller provides support for auto import of foreign configuration without requiring user interactions. The auto import can be enabled or disabled. If enabled, the PERC controller can auto import any foreign configuration detected without manual intervention. If disabled the PERC does not auto import any foreign configuration.

You must have Login and Server Control privilege to import foreign configurations.

This task is not supported on PERC hardware controllers running in HBA mode.

NOTE: It is not recommended to remove an external enclosure cable while the operating system is running on the system. Removing the cable could result in a foreign configuration when the connection is re-established.

You can manage foreign configurations in the following cases:

- All the physical disks in a configuration are removed and re-inserted.
- Some of the physical disks in a configuration are removed and re-inserted.
- All the physical disks in a virtual disk are removed, but at different times, and then re-inserted.
- The physical disks in a non-redundant virtual disk are removed.

The following constraints apply to the physical disks that are considered for import:



- The drive state of a physical disk can change from the time the foreign configuration is scanned to when the actual import
 occurs. The foreign import occurs only on drives that are in the Unconfigured Good state.
- Drives in the failed or offline state cannot be imported.
- The firmware does not allow you to import more than eight foreign configurations.

Importing foreign configuration using web interface

NOTE: If there is an incomplete foreign disk configuration in the system, then the state of one or more existing online virtual disks is also displayed as foreign.

(i) NOTE: Importing foreign configuration for BOSS controller is not supported.

To import foreign configuration:

- 1. In the iDRAC web interface, go to Configuration > Storage Configuration;.
- 2. From the Controller drop-down menu, select the controller you want to import the foreign configuration to.
- 3. Click Import under the Foreign Configuration and then click Apply.

Importing foreign configuration using RACADM

To import foreign configuration:

racadm storage importconfig:<Controller FQDD>

For more information, see the iDRAC RACADM Command Line Reference Guide available at dell.com/idracmanuals.

Clearing foreign configuration

After moving a physical disk from one controller to another, you may find that the physical disk contains all or some portion of a virtual disk (foreign configuration). You can identify whether a previously used physical disk contains a foreign configuration (virtual disk) by checking the physical disk state. If the physical disk state is Foreign, then the physical disk contains all or some portion of a virtual disk. You can clear or erase the virtual disk information from the newly attached physical disks.

The Clear Foreign Configuration operation permanently erases all data residing on the physical disks that are added to the controller. If more than one foreign virtual disk is present, all the configurations are erased. You may prefer to import the virtual disk rather than destroy the data. An initialization must be performed to remove foreign data. If you have an incomplete foreign configuration which cannot be imported, you can use the Clearing Foreign Configuration option to erase the foreign data on the physical disks.

Clearing foreign configuration using web interface

To clear the foreign configuration:

- 1. In the iDRAC web interface, go to Configuration > Storage Configuration > Controller Configuration. The Controller Configuration page is displayed.
- 2. From the **Controller** drop-down menu, select the controller for which you want to clear the foreign configuration.

(i) NOTE: To clear foreign configuration on BOSS controllers, click "Reset Configuration".

3. Click Clear Configuration.

4. Click Apply

Based on the selected operation mode, the virtual disks residing on the physical disk is erased.



Clearing foreign configuration using RACADM

To clear foreign configuration:

racadm storage clearconfig:<Controller FQDD>

For more information, see the *iDRAC RACADM Command Line Reference Guide* available at **dell.com/idracmanuals**.

Resetting controller configuration

You can reset the configuration for a controller. This operation deletes virtual disk drives and unassigns all hot spares on the controller. It does not erase any data other than removing the disks from the configuration. Reset configuration also does not remove any foreign configurations. The real-time support of this feature is available only in PERC 9.1 firmware. Reset configuration does not erase any data. You may recreate the exact same configuration without an initialize operation which may result in the data being recovered. You must have server control privilege.

NOTE: Resetting the controller configuration does not remove a foreign configuration. To remove a foreign configuration, perform clear configuration operation.

Resetting controller configuration using web interface

To reset the controller configuration:

- 1. In the iDRAC Web interface, go to **Storage** > **Overview** > **Controllers**.
- 2. From the Actions, select Reset Configuration for one or more controllers.
- 3. For each controller, from the Apply Operation Mode drop-down menu, select when you want to apply the settings.
- Click Apply.
 Based on the selected operation mode, the settings are applied.

Resetting controller configuration using RACADM

To reset the controller configuration:

racadm storage resetconfig:<Controller FQDD>

For more information, see the iDRAC RACADM Command Line Reference Guide available at dell.com/idracmanuals.

Switching the controller mode

On PERC 9.1 controllers, you can change the personality of the controller by switching the mode from RAID to HBA. The controller operates similar to an HBA controller where the drivers are passed through the operating system. The controller mode change is a staged operation and does not occur in real time.

PERC 10 and later controllers supports Enhanced HBA mode, replacing HBA from the current controller mode options. However, PERC 9 still continues to support the HBA mode.

() NOTE:

- Enhanced HBA supports Non-RAID PDs and all RAID level VDs.
- It only supports creation of RAID0, RAID1, and RAID10 VDs.
- Enhanced HBA is not supported on PERC 11.

Enhanced HBA mode provides the following features:

- Create virtual disks with RAID level 0, 1, or 10.
- Present non-RAID disks to host.
- Configure a default cache policy for virtual disks as write-back with read ahead.
- Configure virtual disks and non-RAID disks as valid boot devices.
- Automatically convert all unconfigured disks to non-RAID:
 - On system boot



- On controller reset
- When unconfigured disks are hot-inserted

NOTE: Creating or importing RAID 5, 6, 50, or 60 virtual disks is not supported. Also, in enhanced HBA mode, non-RAID disks are enumerated first in ascending order, while RAID volumes are enumerated in descending order.

Before you change the mode of the controller from RAID to HBA, ensure that:

- The RAID controller supports the controller mode change. The option to change the controller mode is not available on controllers where the RAID personality requires a license.
- All virtual disks must be deleted or removed.
- Hot spares must be deleted or removed.
- Foreign configurations must be deleted or cleared.
- All physical disks that are in a failed state, must be removed or the pinned cache needs to be cleared.
- Any local security key that is associated with SEDs must be deleted.
- The controller must not have preserved cache.
- You have server control privileges to switch the controller mode.
- () NOTE: Ensure that you back up the foreign configuration, security key, virtual disks, and hot spares before you switch the mode as the data is deleted.
- () NOTE: Ensure that a CMC license (not applicable for MX platforms) is available for PERC FD33xS and FD33xD storage sleds before you change the controller mode. For more information on CMC license for the storage sleds, see the *Dell Chassis Management Controller Version 1.2 for PowerEdge FX2/FX2s User's Guide* available at dell.com/cmcmanuals .

Exceptions while switching the controller mode

The following list provides the exceptions while setting the controller mode using the iDRAC interfaces such as web interface, RACADM, and WSMan:

- If the PERC controller is in RAID mode, you must clear any virtual disks, hot spares, foreign configurations, controller keys, or preserved cache before changing it to HBA mode.
- You cannot configure other RAID operations while setting the controller mode. For example, if the PERC is in RAID mode and you set the pending value of the PERC to HBA mode, and you try to set the BGI attribute, the pending value is not initiated.
- When you switch the PERC controller from HBA to RAID mode, the drives remain in Non-RAID state and are not automatically set to Ready state. Additionally, the **RAIDEnhancedAutoImportForeignConfig** attribute is automatically set to **Enabled**.

The following list provides the exceptions while setting the controller mode using the Server Configuration Profile feature using the WSMan or RACADM interface:

- Server Configuration Profile feature allows you to configure multiple RAID operations along with setting the controller mode. For example, if the PERC controller is in HBA mode, you can edit the export Server Configuration Profile (SCP) to change the controller mode to RAID, convert drives to ready and create a virtual disk.
- While changing the mode from RAID to HBA, the **RAIDaction pseudo** attribute is set to update (default behavior). The attribute runs and creates a virtual disk which fails. The controller mode is changed, however, the job is completed with errors. To avoid this issue, you must comment out the RAIDaction attribute in the SCP file.
- When the PERC controller is in HBA mode, if you run import preview on export SCP which is edited to change controller mode to RAID, and try creating a VD, the virtual disk creation fails. Import preview does not support validating stacking RAID operations with changing controller mode.

Switching the controller mode using the iDRAC web interface

To switch the controller mode, perform the following steps:

- 1. In the iDRAC web interface, click Storage > Overview > Controllers.
- On the Controllers page, click Action > Edit.The Current Value column displays the current setting of the controller.
- **3.** From the drop-down menu, select the controller mode you want to switch to, and click **At Next Reboot**. Reboot the system for the change to take effect.



Switching the controller mode using RACADM

To switch the controller mode using RACADM, run the following commands.

To view the current mode of the controller:

\$ racadm get Storage.Controller.1.RequestedControllerMode[key=<Controller_FQDD>]

The following output is displayed:

RequestedControllerMode = NONE

• To set the controller mode as HBA:

\$ racadm set Storage.Controller.1.RequestedControllerMode HBA [Key=<Controller FQDD>]

• To create a job and apply changes:

\$ racadm jobqueue create <Controller Instance ID> -s TIME_NOW -r pwrcycle

For more information, see the *iDRAC RACADM Command Line Interface Reference Guide* available at **dell.com/idracmanuals**.

12 Gbps SAS HBA adapter operations

Dell PowerEdge servers must have an operating system installed and the appropriate device driver to be loaded in order for Dell HBAs to operate. Following POST, the HBA ports will be disabled. The HBA device driver is responsible for resetting the HBA and enabling its ports connected to storage devices. Without an operating system, the driver will not be loaded, and there is no guarantee that iDRAC will be able to display storage devices connected to Dell HBAs.

The non-RAID controllers are the HBAs that do not have few RAID capabilities. They do not support virtual disks.

14G iDRAC interface supports 12 Gbps SAS HBA controller, HBA330 (integrated and adapter) controllers, HBA330 MMZ, and HBA330 MX adapters.

AMD platforms support HBA355i front and HBA355i Adapter controllers.

You can perform the following for non-RAID controllers:

- View controller, physical disks, and enclosure properties as applicable for the non-RAID controller. Also, view EMM, fan, power supply unit, and temperature probe properties associated with the enclosure. The properties are displayed based on the type of controller.
- View software and hardware inventory information.
- Update firmware for enclosures behind the 12 Gbps SAS HBA controller (staged)
- Monitor the polling or polling frequency for physical disk SMART trip status when there is change detected
- Monitor the physical disks hot plug or hot removal status
- Blink or unblink LEDs

() NOTE:

- There is limited support for tape drives when they are connected behind 12gbps SAS or HBA355e.
- Even though LED is not available for tape drive, blink/unblink option can be successful.

() NOTE:

- Enable Collect System Inventory On Reboot (CSIOR) operation before inventorying or monitoring the non-RAID controllers.
- Real-time monitoring for SMART enabled drives and SES enclosure sensors is only done for the 12 Gbps SAS HBA controllers and HBA330 internal controllers.

(i) NOTE: Detection of failed drives behind SAS HBA controllers is not supported.

Monitoring predictive failure analysis on drives

Storage management supports Self Monitoring Analysis and Reporting Technology (SMART) on physical disks that are SMARTenabled.



SMART performs predictive failure analysis on each disk and sends alerts if a disk failure is predicted. The controllers check physical disks for failure predictions and, if found, pass this information to iDRAC. iDRAC immediately logs an alert.

Controller operations in non-RAID mode or HBA mode

If the controller is in non-RAID mode (HBA mode), then:

- Virtual disks or hot spares are not available.
- Security state of the controller is disabled.
- All physical disks are in non-RAID mode.

You can perform the following operations if the controller is in non-RAID mode:

- Blink/unblink the physical disk.
- Configure all properties including the following:
- Load balanced mode
- Check consistency mode
- Patrol read mode
- Copyback mode
- Controller boot mode
- Enhanced auto import foreign configuration
- Rebuild rate
- Check consistency rate
- Reconstruct rate
- o BGI rate
- Enclosure or backplane mode
- Patrol read unconfigured areas
- View all properties that are applicable to a RAID controller expect for virtual disks.
- Clear foreign configuration

(i) NOTE: If an operation is not supported in non-RAID mode, an error message is displayed.

You cannot monitor the enclosure temperature probes, fans, and power supplies when the controller is in non-RAID mode.

Running RAID configuration jobs on multiple storage controllers

While performing operations on more than two storage controllers from any supported iDRAC interface, make sure to:

- Run the jobs on each controller individually. Wait for each job to complete before starting the configuration and job creation on the next controller.
- Schedule multiple jobs to run at a later time using the scheduling options.

Manage Preserved cache

The Managed Preserved Cache feature is a controller option which provides the user an option to discard the controller cache data. In the write-back policy, data is written to the cache before being written to the physical disk. If the virtual disk goes offline or is deleted for any reason, the data in the cache gets deleted.

The PREC Controller preserves the data written on the preserved or dirty cache in an event of power failure or cable disconnect until you recover the virtual disk or clear the cache.

The status of the controller is affected by the preserved cache. The controller status is displayed as degraded if the controller has preserved cache. Discard the preserved cache is possible only if all of the following conditions are met:

- The controller does not have any foreign configuration.
- The controller does not have any offline or missing virtual disks.
- Cables to any virtual disk are not disconnected.



Managing PCIe SSDs

Peripheral Component Interconnect Express (PCIe) solid-state device (SSD) is a high-performance storage device designed for solutions requiring low latency, high Input Output Operations per Second (IOPS), and enterprise class storage reliability and serviceability. The PCIe SSD is designed based on Single Level Cell (SLC) and Multi-Level Cell (MLC) NAND flash technology with a high-speed PCIe 2.0, PCIe 3.0, or PCIe 4.0 compliant interface. In 14th generation of PowerEdge servers, we have three different ways to connect SSDs. You can use an extender to connect the SSDs via backplane, directly connect the SSDs from backplane to mother board using slimline cable without extender, and use HHHL (Add-In) card which sits on the motherboard.

() NOTE:

- 14th generation of PowerEdge servers are supporting Industry standard NVMe-MI specification based NVMe SSDs
- PERC 11 supports PCIe SSD/NVMe devices behind PERC inventory monitoring and configuration.

Using iDRAC interfaces, you can view and configure NVMe PCIe SSDs.

The key features of PCIe SSD are:

- Hot plug capability
- High-performance device

In few of the 14th generation of PowerEdge servers, up to 32 NVMe SSDs are supported.

You can perform the following operations for PCIe SSDs:

- Inventory and remotely monitor the health of PCIe SSDs in the server
- Prepare to remove the PCIe SSD
- Securely erase the data
- Blink or unblink the device LED (Identify the device)
- You can perform the following operations for HHHL SSDs:
- Inventory and real-time monitoring of the HHHL SSD in the server
- Failed card reporting and logging in iDRAC and OMSS
- Securely erasing the data and removing the card
- TTY logs reporting
- You can perform the following operations for SSDs:
- Drive status reporting such as Online, Failed, and Offline
- **NOTE:** Hot plug capability, prepare to remove, and blink or unblink the device LED is not applicable for HHHL PCIe SSD devices.
- **NOTE:** When NVMe devices are controlled behind S140, prepare to remove and cryptographic erase operations are not supported, blink and unblink are supported.

Inventorying and monitoring PCIe SSDs

The following inventory and monitoring information is available for PCle SSDs:

- Hardware information:
- PCIe SSD Extender card
- PCIe SSD Backplane

If the system has a dedicated PCIe backplane, two FQDDs are displayed. One FQDD is for regular drives and the other is for SSDs. If the backplane is shared (universal), only one FQDD is displayed. In case the SSDs are directly attached, the controller FQDD reports as CPU.1, indicating that the SSD is directly attached to the CPU.

• Software inventory includes only the firmware version for the PCIe SSD.

Inventorying and monitoring PCIe SSDs using web interface

To inventory and monitor PCle SSD devices, in the iDRAC web interface, go to **Storage** > **Overview** > **Physical Disks**. The **Properties** page is displayed. For PCle SSDs, the **Name** column displays **PCle SSD**. Expand to view the properties.



Inventorying and monitoring PCIe SSDs using RACADM

Use the racadm storage get controllers: <PcieSSD controller FQDD> command to inventory and monitor PCle SSDs.

To view all PCIe SSD drives:

racadm storage get pdisks

To view PCIe extender cards:

racadm storage get controllers

To view PCIe SSD backplane information:

racadm storage get enclosures

(i) NOTE: For all the mentioned commands, PERC devices are also displayed.

For more information, see the *iDRAC RACADM Command Line Reference Guide* available at **dell.com/idracmanuals**.

Preparing to remove PCIe SSD

() NOTE: This operation is not supported when:

- PCIe SSD is configured using the S140 controller.
- NVMe device is behind PERC 11.

PCIe SSDs support orderly hot swap allowing you to add or remove a device without halting or rebooting the system in which the devices are installed. To prevent data loss, you must use the Prepare to Remove operation before physically removing a device.

Orderly hot swap is supported only when PCIe SSDs are installed in a supported system running a supported operating system. To ensure that you have the correct configuration for your PCIe SSD, see the system-specific owner's manual.

The Prepare to Remove operation is not supported for PCIe SSDs on the VMware vSphere (ESXi) systems and HHHL PCIe SSD devices.

NOTE: Prepare to Remove operation is supported on systems with ESXi 6.0 with iDRAC Service Module version 2.1 or higher.

The Prepare to Remove operation can be performed in real-time using iDRAC Service Module.

The Prepare to Remove operation stops any background activity and any ongoing I/O activity so that device can be removed safely. It causes the status LEDs on the device to blink. You can safely remove the device from the system under the following conditions after you initiate the Prepare to Remove operation:

- The PCIe SSD is blinking the safe to remove LED pattern (blinks amber).
- The PCIe SSD is no longer accessible by the system.

Before preparing the PCle SSD for removal, ensure that:

- iDRAC Service Module is installed.
- Lifecycle Controller is enabled.
- You have Server Control and Login privileges.

Preparing to remove PCIe SSD using web interface

To prepare the PCIe SSD for removal:

- In the iDRAC Web interface, go to Storage > Overview > Physical Disks. The Setup Physical Disk page is displayed.
- 2. From the Controller drop-down menu, select the extender to view the associated PCIe SSDs.
- 3. From the drop-down menus, select **Prepare to Remove** for one or more PCIe SSDs.

If you have selected **Prepare to Remove** and you want to view the other options in the drop-down menu, then select **Action** and then click the drop-down menu to view the other options.



(i) NOTE: Ensure that iSM is installed and running to perform the preparetoremove operation.

From the Apply Operation Mode drop-down menu, select Apply Now to apply the actions immediately.
 If there are jobs to be completed, then this option is grayed-out.

(i) NOTE: For PCIe SSD devices, only the Apply Now option is available. This operation is not supported in staged mode.

5. Click Apply.

If the job is not created, a message indicating that the job creation was not successful is displayed. Also, the message ID and the recommended response action is displayed.

If the job is created successfully, a message indicating that the job ID is created for the selected controller is displayed. Click **Job Queue** to view the progress of the job in the **Job Queue** page.

If pending operation is not created, an error message is displayed. If pending operation is successful and job creation is not successful, then an error message is displayed.

Preparing to remove PCIe SSD using RACADM

To prepare the PCIeSSD drive for removal:

racadm storage preparetoremove:<PCIeSSD FQDD>

To create the target job after executing preparetoremove command:

racadm jobqueue create <PCIe SSD FQDD> -s TIME_NOW --realtime

To query the job ID returned:

racadm jobqueue view -i <job ID>

For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Erasing PCIe SSD device data

(i) NOTE: This operation is not supported when PCIe SSD is configured using the SWRAID controller.

Cryptographic Erase permanently erases all data present on the disk. Performing a Cryptographic Erase on an PCIe SSD overwrites all blocks and results in permanent loss of all data on the PCIe SSD. During Cryptographic Erase, the host is unable to access the PCIe SSD. The changes are applied after system reboot.

If the system reboots or experiences a power loss during cryptographic erase, the operation is canceled. You must reboot the system and restart the process.

Before erasing PCIe SSD device data, make sure that:

- Lifecycle Controller is enabled.
- You have Server Control and Login privileges.

() NOTE:

- Erasing PCIe SSDs can only be performed as a staged operation.
- After the drive is erased, it displays in the operating system as online but it is not initialized. You must initialize and format the drive before using it again.
- After you hot-plug a PCIe SSD, it may take several seconds to appear on the web interface.

Erasing PCIe SSD device data using web interface

To erase the data on the PCIe SSD device:

 In the iDRAC Web interface, go to Storage > Overview > Physical Disks. The Physical Disk page is displayed.



- 2. From the **Controller** drop-down menu, select the controller to view the associated PCle SSDs.
- 3. From the drop-down menus, select **Cryptographic Erase** for one or more PCIe SSDs.

If you have selected **Cryptographic Erase** and you want to view the other options in the drop-down menu, then select **Action** and then click the drop-down menu to view the other options.

- 4. From the Apply Operation Mode drop-down menu, select one of the following options:
 - At Next Reboot Select this option to apply the actions during the next system reboot.
 - At Scheduled Time Select this option to apply the actions at a scheduled day and time:
 - Start Time and End Time Click the calendar icons and select the days. From the drop-down menus, select the time. The action is applied between the start time and end time.
 - From the drop-down menu, select the type of reboot:
 No Reboot (Manually Reboot System)

 - Graceful Shutdown
 - Force Shutdown
 - Power Cycle System (cold boot)

5. Click Apply.

If the job is not created, a message indicating that the job creation was not successful is displayed. Also, the message ID and the recommended response action is displayed.

If the job is created successfully, a message indicating that the job ID is created for the selected controller is displayed. Click **Job Queue** to view the progress of the job in the Job Queue page.

If pending operation is not created, an error message is displayed. If pending operation is successful and job creation is not successful, then an error message is displayed.

Erasing PCIe SSD device data using RACADM

To securely erase a PCIe SSD device:

```
racadm storage secureerase:<PCIeSSD FQDD>
```

To create the target job after executing the secureerase command:

racadm jobqueue create <PCIe SSD FQDD> -s TIME NOW -e <start time>

To query the job ID returned:

racadm jobqueue view -i <job ID>

For more information, see the iDRAC RACADM Command Line Reference Guide available at dell.com/idracmanuals.

Managing enclosures or backplanes

You can perform the following for enclosures or backplanes:

- View properties
- Configure universal mode or split mode
- View slot information (universal or shared)
- Set SGPIO mode
- Set Asset Tag
- Asset Name

Configuring backplane mode

The Dell 14th generation PowerEdge servers supports a new internal storage topology, where two storage controllers (PERCs) can be connected to a set of internal drives through a single expander. This configuration is used for high performance mode with no failover or High Availability (HA) functionality. The expander splits the internal drive array between the two storage



controllers. In this mode, virtual disk creation only displays the drives connected to a particular controller. There are no licensing requirements for this feature. This feature is supported only on a few systems.

Backplane supports the following modes:

- Unified mode This is the default mode. The primary PERC controller has access to all the drives connected to the backplane even if a second PERC controller is installed.
- Split mode One controller has access to the first 12 drives and the second controller has access to the last 12 drives. The drives connected to the first controller are numbered 0-11 while the drives connected to the second controller are numbered 12-23.
- Split mode 4:20 One controller has access to the first 4 drives and the second controller has access to the last 20 drives. The drives connected to the first controller are numbered 0-3 while the drives connected to the second controller are numbered 4-23.
- Split mode 8:16 One controller has access to the first 8 drives and the second controller has access to the last 16 drives. The drives connected to the first controller are numbered 0-7 while the drives connected to the second controller are numbered 8-23.
- Split mode 16:8 One controller has access to the first 16 drives and the second controller has access to the last 8 drives. The drives connected to the first controller are numbered 0-15 while the drives connected to the second controller are numbered 16-23.
- Split mode 20:4 One controller has access to the first 20 drives and the second controller has access to the last 4 drives. The drives connected to the first controller are numbered 0-19 while the drives connected to the second controller are numbered 20-23.
- Split mode 6:6:6:6 4 blades installed in one chassis and each blade have 6 drives assigned. This mode is only supported on PowerEdge C series blades.
- Information Not Available Controller information is not available.

iDRAC allows the split mode setting if the expander has the capability to support the configuration. Ensure that you enable this mode prior to installing the second controller. iDRAC performs a check for expander capability prior to allowing this mode to be configured and does not check whether the second PERC controller is present.

() NOTE: Cable errors (or other errors) may appear if you put the backplane into split mode with only one PERC attached, or if you put the backplane into Unified Mode with two PERCs attached.

To modify the setting, you must have Server Control privilege.

If any other RAID operations are in pending state or any RAID job is scheduled, you cannot change the backplane mode. Similarly, if this setting is pending, you cannot schedule other RAID jobs.

() NOTE:

- Warning messages are displayed when the setting is being changed as there is a possibility of data loss.
- LC Wipe or iDRAC reset operations do not change the expander setting for this mode.
- This operation is supported only in real-time and not staged.
- You can change the backplane configuration multiple times.
- The backplane splitting operation can cause data loss or foreign configuration if the drive association changes from one controller to another controller.
- During the backplane splitting operation, the RAID configuration may be impacted depending on the drive association.

Any change in this setting only takes effect after a system power reset. If you change from Split mode to Unified, an error message is displayed on the next boot as the second controller does not see any drives. Also, the first controller will see a foreign configuration. If you ignore the error, the existing virtual disks are lost.

Configuring backplane mode using web interface

To configure backplane mode using iDRAC web interface:

- 1. In the iDRAC web interface, go to Configuration > Storage Configuration > Enclosures Configuration.
- 2. From the **Controller** menu, select the controller to configure its associated enclosures.
- 3. From the Action drop-down menu, select Edit Enclosure Mode.
- The Edit Enclosure Mode page is displayed.
- 4. In the **Current Value** column, select the required enclosure mode for the backplane or enclosure. The options are:
 - Unified Mode
 - Split Mode
 - Split Mode 4:20



- Split Mode 8:16
- Split Mode 16:8
- Split Mode 20:4

NOTE: For C6420, the available modes are: Split Mode and Split Mode-6:6:6:6. Few values may be only supported on certain platforms.

For R740xd and R940, power cycle of the server is needed to apply the new backplane zone and for C6420, A/C cycle (of the blade chassis) to apply the new backplane zone.

- Click Add to Pending Operations. A job ID is created.
- 6. Click Apply Now.
- 7. Go to the **Job Queue** page and verify that it displays the status as Completed for the job.
- 8. Power cycle the system for the setting to take effect.

Configuring enclosure using RACADM

To configure the enclosure or backplane, use the set command with the objects in **BackplaneMode**.

For example, to set the BackplaneMode attribute to split mode:

1. Run the following command to view the current backplane mode:

racadm get storage.enclosure.1.backplanecurrentmode

The output is:

BackplaneCurrentMode=UnifiedMode

2. Run the following command to view the requested mode:

racadm get storage.enclosure.1.backplanerequestedmode

The output is:

BackplaneRequestedMode=None

3. Run the following command to set the requested backplane mode to split mode:

racadm set storage.enclosure.1.backplanerequestedmode "splitmode"

The message is displayed indicating that the command is successful.

4. Run the following command to verify if the **backplanerequestedmode** attribute is set to split mode:

racadm get storage.enclosure.1.backplanerequestedmode

The output is:

BackplaneRequestedMode=None (Pending=SplitMode)

- 5. Run storage get controllers command and note down the controller instance ID.
- 6. Run the following command to create a job:

racadm jobqueue create <controller instance ID> -s TIME_NOW --realtime

A job ID is returned.

7. Run the following command to query the job status:

racadm jobqueue view -i JID_xxxxxxx

264 Managing storage devices



where, JID XXXXXXX is the job ID from step 6.

The status is displayed as Pending.

Continue to query the job ID until you view the Completed status (this process may take up to three minutes).

 $\textbf{8.} \ \ \mbox{Run the following command to view the backplanerequested mode attribute value:}$

 $\verb| racadm get storage.enclosure.1.backplanerequestedmode| \\$

The output is:

BackplaneRequestedMode=SplitMode

9. Run the following command to cold reboot the server:

racadm serveraction powercycle

10. After the system completes POST and CSIOR, type the following command to verify the backplanerequestedmode:

racadm get storage.enclosure.1.backplanerequestedmode

The output is:

BackplaneRequestedMode=None

11. Run the following to verify is the backplane mode is set to split mode:

racadm get storage.enclosure.1.backplanecurrentmode

The output is:

BackplaneCurrentMode=SplitMode

12. Run the following command and verify that only 0–11 drives are displayed:

racadm storage get pdisks

For more information about the RACADM commands, see the *iDRAC RACADM Command Line Interface Reference Guide* available at **dell.com/idracmanuals**.

Viewing universal slots

Some 14th generation PowerEdge server backplanes supports both SAS/SATA and PCIe SSD drives in the same slot. These slots are called universal slots and are wired to the primary storage controller (PERC) and either PCIe extender card or direct connect manager by CPU backplanes supports both SAS/SATA and PCIe SSD drives in the same slot. The backplane firmware provides information about the slots that support this feature. The backplane supports SAS/SATA disks or PCIe SSDs. Typically, the four higher number slots are universal. For example, in a universal backplane that supports 24 slots, slots 0-19 support only SAS/SATA disks while slots 20-23 support either SAS/SATA or PCIe SSD.

The roll-up health status for the enclosure provides the combined health status for all the drives in the enclosure. The enclosure link on the **Topology** page displays the entire enclosure information irrespective of which controller it is associated with. Because two storage controllers (PERC and PCle extender) can be connected to the same backplane, only the backplane associated with the PERC controller is displayed in **System Inventory** page.

In the Storage > Enclosures > Properties page, the Physical Disks Overview section displays the following:

- **Slot Empty** If a slot is empty.
- PCIe Capable If there are no PCIe capable slots, this column is not displayed.
- Bus Protocol If it is a universal backplane with PCIe SSD installed in one of the slots, this column displays PCIe.
- Hotspare This column is not applicable for PCIe SSD.

() NOTE: Hot swapping is supported for universal slots. If you want to remove a PCIe SSD drive and swap it with a SAS/ SATA drive, ensure that you first complete the PrepareToRemove task for the PCIe SSD drive. If you do not perform this task, the host operating system may have issues such as a blue screen, kernel panic, and so on.



Setting SGPIO mode

The storage controller can connect to the backplane in I2C mode (default setting for Dell backplanes) or Serial General Purpose Input/Output (SGPIO) mode. This connection is required for blinking LEDs on the drives. Dell PERC controllers and backplane support both these modes. To support certain channel adapters, the backplane mode must be changed SGPIO mode.

The SGPIO mode is only supported for passive backplanes. It is not supported for expander-based backplanes or passive backplanes in downstream mode. Backplane firmware provides information on capability, current state, and requested state.

After LC wipe operation or iDRAC reset to default, the SGPIO mode is reset to disabled state. It compares the iDRAC setting with the backplane setting. If the backplane is set to SGPIO mode, iDRAC changes its setting to match the backplane setting.

Server power cycle is required for any change in setting to take effect.

You must have Server Control privilege to modify this setting.

(i) NOTE: You cannot set the SGPIO mode using iDRAC Web interface.

Setting SGPIO mode using RACADM

To configure the SGPIO mode, use the set command with the objects in the SGPIOMode group.

If it is set to disabled, it is I2C mode. If enabled, it is set to SGPIO mode.

For more information, see the *iDRAC RACADM Command Line Interface Reference Guide* available at **dell.com/idracmanuals**.

Set Enclosure Asset Tag

Set Enclosure Asset Tag allows you to configure Asset Tag of a storage enclosure.

User can change the Asset Tag property of the enclosure to identify enclosures. These fields are checked for invalid values and an error is displayed if an invalid value is entered. These fields are part of the enclosure firmware; the data initially shown are the values saved in the firmware.

(i) NOTE: Asset Tag has a character limit of 10 that includes the null character.

(i) NOTE: These operations are not supported on internal enclosures.

Set Enclosure Asset Name

Set Enclosure Asset Name allows the user to configure the Asset Name of a storage enclosure.

The user can change the Asset Name property of the enclosure to identify enclosures easily. These fields are checked for invalid values and an error is displayed if an invalid value is entered. These fields are part of the enclosure firmware; the data initially shown are the values saved in the firmware.

(i) NOTE: Asset Name has a character limit of 32 that includes the null character.

(i) NOTE: These operations are not supported on internal enclosures.

Choosing operation mode to apply settings

While creating and managing virtual disks, setting up physical disks, controllers, and enclosures or resetting controllers, before you apply the various settings, you must select the operation mode. That is, specify when you want to apply the settings:

- Immediately
- During the next system reboot
- At a scheduled time
- As a pending operation to be applied as a batch as part of a single job.



Choosing operation mode using web interface

To select the operation mode to apply the settings:

- 1. You can select the operation mode on when you are on any of the following pages:
 - Storage > Physical Disks .
 - Storage > Virtual Disks
 - Storage > Controllers
 - Storage > Enclosures
- 2. Select one of the following from the Apply Operation Mode drop-down menu:
 - **Apply Now** Select this option to apply the settings immediately. This option is available for PERC 9 controllers only. If there are jobs to be completed, then this option is grayed-out. This job take at least 2 minutes to complete.
 - At Next Reboot Select this option to apply the settings during the next system reboot.
 - At Scheduled Time Select this option to apply the settings at a scheduled day and time:
 - **Start Time** and **End Time** Click the calendar icons and select the days. From the drop-down menus, select the time. The settings are applied between the start time and end time.
 - \circ $\,$ From the drop-down menu, select the type of reboot:
 - No Reboot (Manually Reboot System)
 - Graceful Shutdown
 - Force Shutdown
 - Power Cycle System (cold boot)
 - Add to Pending Operations Select this option to create a pending operation to apply the settings. You can view all pending operations for a controller in the Storage > Overview > Pending Operations page.
 - (i) NOTE:
 - The Add to Pending Operations option in not applicable for the Pending Operations page and for PCIe SSDs in the Physical Disks > Setup page.
 - Only the **Apply Now** option is available on the **Enclosure Setup** page.
- 3. Click Apply.

Based on the operation mode selected, the settings are applied.

Choosing operation mode using RACADM

To select the operation mode, use the jobqueue command.

For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Viewing and applying pending operations

You can view and commit all pending operations for the storage controller. All the settings are either applied at once, during the next reboot, or at a scheduled time based on the selected options. You can delete all the pending operations for a controller. You cannot delete individual pending operations.

Pending Operations are created on the selected components (controllers, enclosures, physical disks, and virtual disks).

Configuration jobs are created only on controller. In case of PCIe SSD, job is created on PCIe SSD disk and not on the PCIe Extender.

Viewing, applying, or deleting pending operations using web interface

- In the iDRAC web interface, go to Storage > Overview > Pending Operations. The Pending Operations page is displayed.
- 2. From the **Component** drop-down menu, select the controller for which you want to view, commit, or delete the pending operations.

The list of pending operations is displayed for the selected controller.



() NOTE:

- Pending operations are created for import foreign configuration, clear foreign configuration, security key operations, and encrypt virtual disks. But, they are not displayed in the **Pending Operations** page and in the Pending Operations pop-up message.
- Jobs for PCIe SSD cannot be created from the **Pending Operations** page
- **3.** To delete the pending operations for the selected controller, click **Delete All Pending Operations**.
- 4. From the drop-down menu, select one of the following and click Apply to commit the pending operations:
 - **Apply Now** Select this option to commit all the operations immediately. This option is available for PERC 9 controllers with the latest firmware versions.
 - At Next Reboot Select this option to commit all the operations during the next system reboot.
 - At Scheduled Time Select this option to commit the operations at a scheduled day and time.
 - **Start Time** and **End Time** Click the calendar icons and select the days. From the drop-down menus, select the time. The action is applied between the start time and end time.
 - From the drop-down menu, select the type of reboot:
 - No Reboot (Manually Reboot System)
 - Graceful Shutdown
 - Force Shutdown
 - Power Cycle System (cold boot)
- 5. If the commit job is not created, a message indicating that the job creation was not successful is displayed. Also, the message ID and the recommended response action are displayed.
- 6. If the commit job is created successfully, a message indicating that the job ID is created for the selected controller is displayed. Click **Job Queue** to view the progress of the job in the **Job Queue** page.

If the clear foreign configuration, import foreign configuration, security key operations, or encrypt virtual disk operations are in pending state, and if these are the only operations pending, then you cannot create a job from the **Pending Operations** page. You must perform any other storage configuration operation or use RACADM or WSMan to create the required configuration job on the required controller.

You cannot view or clear pending operations for PCle SSDs in the **Pending Operations** page. Use the racadm command to clear the pending operations for PCle SSDs.

Viewing and applying pending operations using RACADM

To apply pending operations, use the **jobqueue** command.

For more information, see the *iDRAC RACADM Command Line Reference Guide* available at **dell.com/idracmanuals**.

Storage devices — apply operation scenarios

Case 1: selected an apply operation (apply now, at next reboot, or at scheduled time) and there are no existing pending operations

If you have selected **Apply Now**, **At Next Reboot**, or **At Scheduled Time** and then clicked **Apply**, first the pending operation is created for the selected storage configuration operation.

- If the pending operation is successful and there are no prior existing pending operations, then the job is created. If the job is created successfully, a message indicating that the job ID is created for the selected device is displayed. Click **Job Queue** to view the progress of the job in the **Job Queue** page. If the job is not created, a message indicating that the job creation was not successful is displayed. Also, the message ID and the recommended response action are displayed.
- If the pending operation creation is unsuccessful and there are no prior existing pending operations, an error message with ID and recommended response action is displayed.

Case 2: selected an apply operation (apply now, at next reboot, or at scheduled time) and there are existing pending operations

If you have selected **Apply Now**, **At Next Reboot**, or **At Scheduled Time** and then clicked **Apply**, first the pending operation is created for the selected storage configuration operation.

• If the pending operation is created successfully and if there are existing pending operations, then a message is displayed.



- Click the View Pending Operations link to view the pending operations for the device.
- Click Create Job to create job for the selected device. If the job is created successfully, a message indicating that the job ID is created for the selected device is displayed. Click Job Queue to view the progress of the job in the Job Queue page. If the job is not created, a message indicating that the job creation was not successful is displayed. Also, the message ID and the recommended response action is displayed.
- Click Cancel to not create the job and remain on the page to perform more storage configuration operations.
- If the pending operation is not created successfully and if there are existing pending operations, then an error message is displayed.
- \circ $\;$ Click $\mbox{Pending Operations}$ to view the pending operations for the device.
- Click Create Job For Successful Operations to create the job for the existing pending operations. If the job is created successfully, a message indicating that the job ID is created for the selected device is displayed. Click Job Queue to view the progress of the job in the Job Queue page. If the job is not created, a message indicating that the job creation was not successful is displayed. Also, the message ID and the recommended response action are displayed.

• Click **Cancel** to not create the job and remain on the page to perform more storage configuration operations.

Case 3: selected add to pending operations and there are no existing pending operations

If you have selected **Add to Pending Operations** and then clicked **Apply**, first the pending operation is created for the selected storage configuration operation.

- If the pending operation is created successfully and if there are no existing pending operations, then an information message is displayed:
 - Click **OK** to remain on the page to perform more storage configuration operations.
 - Click **Pending Operations** to view the pending operations for the device. Until the job is created on the selected controller, these pending operations are not applied.
- If the pending operation is not created successfully and if there are no existing pending operations, then an error message is displayed.

Case 4: selected add to pending operations and there are prior existing pending operations

If you have selected **Add to Pending Operations** and then clicked **Apply**, first the pending operation is created for the selected storage configuration operation.

- If the pending operation is created successfully and if there are existing pending operations, then an information message is displayed:
 - $\circ~$ Click OK to remain on the page to perform more storage configuration operations.
 - Click **Pending Operations** to view the pending operations for the device.
- If the pending operation is not created successfully and if there are existing pending operations, then an error message is displayed.
 - $\circ~$ Click OK to remain on the page to perform more storage configuration operations.
 - Click Pending Operations to view the pending operations for the device.

() NOTE:

- At any time, if you do not see the option to create a job on the storage configuration pages, go to Storage Overview > Pending Operations page to view the existing pending operations and to create the job on the required controller.
- Only cases 1 and 2 are applicable for PCIe SSD. You cannot view the pending operations for PCIe SSDs and hence Add
- to Pending Operations option is not available. Use racadm command to clear the pending operations for PCIe SSDs.

Blinking or unblinking component LEDs

You can locate a physical disk, virtual disk drive and PCIe SSDs within an enclosure by blinking one of the Light Emitting Diodes (LEDs) on the disk.

You must have Login privilege to blink or unblink an LED.

The controller must be real-time configuration capable. The real-time support of this feature is available only in PERC 9.1 firmware and later.

(i) NOTE: Blink or unblink is not supported for servers without backplane.

Blinking or unblinking component LEDs using web interface

To blink or unblink a component LED:



- 1. In the iDRAC Web interface, go to any of the following pages as per your requirement:
 - **Storage** > **Overview** > **Physical Disks** > **Status** Displays the identified Physical Disks page where you can blink or unblink the physical disks and PCIe SSDs.
 - Storage > Overview > Virtual Disks > Status- Displays the identified Virtual Disks page where you can blink or unblink the virtual disks.
- 2. If you select the physical disk:
 - Select or deselect all component LEDs Select the **Select/Deselect All** option and click **Blink** to start blinking the component LEDs. Similarly, click **Unblink** to stop blinking the component LEDs.
 - Select or deselect individual component LEDs Select one or more component(s) and click **Blink** to start blinking the selected component LED(s). Similarly, click **Unblink** to stop blinking the component LEDs.
- 3. If you select the virtual disk:
 - Select or deselect all physical disk drives or PCIe SSDs Select the **Select/Deselect All** option and click **Blink** to start blinking all the physical disk drives and the PCIe SSDs. Similarly, click **Unblink** to stop blinking the LEDs.
 - Select or deselect individual physical disk drives or PCIe SSDs Select one or more physical disk drives and click **Blink** to start blinking the LEDs for the physical disk drives or the PCIe SSDs. Similarly, click **Unblink** to stop blinking the LEDs.
- 4. If you are on the Identify Virtual Disk page:
 - Select or deselect all virtual disks Select the **Select/Deselect All** option and click **Blink** to start blinking the LEDs for all the virtual disks. Similarly, click **Unblink** to stop blinking the LEDs.
 - Select or deselect individual virtual disks Select one or more virtual disks and click **Blink** to start blinking the LEDs for the virtual disks. Similarly, click **Unblink** to stop blinking the LEDs.

If the blink or unblink operation is not successful, error messages are displayed.

Blinking or unblinking component LEDs using RACADM

To blink or unblink component LEDs, use the following commands:

racadm storage blink:<PD FQDD, VD FQDD, or PCIe SSD FQDD>

racadm storage unblink:<PD FQDD, VD FQDD, or PCIe SSD FQDD>

For more information, see the iDRAC RACADM Command Line Reference Guide available at dell.com/idracmanuals.

Warm reboot

When warm reboot is performed, following behaviors are observed:

- PERC controllers in iDRAC UI are grayed out immediately after warm reboot. They are available once re-inventory is completed after warm reboot. This is only applicable for PERC controllers and not for NVME/HBA/BOSS.
- Storage files in SupportAssist are empty when PERC controllers are grayed out in GUI.
- LC Logging for PAST event and Critical events are done for PERC during perc reinventory. Rest all LCL for PERC components are suppressed. LCL resumes after PERC re-inventory finishes.
- You cannot start any Real-time job until PERC re-inventory is finished.
- Telemetry data is not collected until PERC re-inventory is finished.
- After the PERC inventory is finished, the behavior is normal.



BIOS Settings

You can view multiple attributes, which are being used for a specific server under the BIOS Settings. You can modify different parameters of each attribute from this BIOS configuration setting. Once you select one attribute, it shows different parameters which are related to that specific attribute. You can modify multiple parameters of an attribute and apply changes before modifying a different attribute. When a user expands a configuration group, attributes are displayed in an alphabetical order.

() NOTE:

- Attribute level help content are dynamically generated.
- The iDRAC Direct USB port is available without host reboot, even when all USB ports are disabled.

Apply

Apply button remains greyed-out until any of the attributes are modified. Once you made changes to an attribute and click **Apply**, it allows you to modify the attribute with required changes. In case, the request fails to set the BIOS attribute, it throws an error with corresponding HTTP Response Status code mapped to SMIL API error or Job Creation error. A message is generated and displayed at that point. For more information, see *Event and Error Message Reference Guide for 14th Generation Dell EMC PowerEdge Servers* available at https://www.dell.com/idracmanuals.

Discard changes

The **Discard Changes** button is greyed-out until any of the attributes are modified. If you click **Discard Changes** button, all the recent changes are discarded and restored with the previous or initial values.

Apply and Reboot

When a user modifies value of an attribute or boot sequence, user is presented with two choices to apply the configuration; **Apply and Reboot** or **Apply on Next Reboot**. In either of the apply options, user is redirected to the job queue page to monitor the progress of that specific job.

A user can view auditing information related to BIOS configuration in the LC logs.

If you click **Apply and Reboot**, it restarts the server immediately to configure all the required changes. In case, the request fails to set the BIOS attributes, it throws an error with corresponding HTTP Response Status code mapped to SMIL API error or Job Creation error. An EEMI message is generated and displayed at that point.

Apply At Next Reboot

When a user modifies value of an attribute or boot sequence, user is presented with two choices to apply the configuration; **Apply and Reboot** or **Apply on Next Reboot**. In either of the apply options, user is redirected to the job queue page to monitor the progress of that specific job.

A user can view auditing information related to BIOS configuration in the LC logs.

If you click **Apply At Next Reboot**, it configures all the required changes on the next restart of the server. You will not experience any immediate modifications based on the recent configuration changes until the next reboot session is taking place successfully. In case, the request fails to set the BIOS attributes, it throws an error with corresponding HTTP Response Status code mapped to SMIL API error or Job Creation error. An EEMI message is generated and displayed at that point.



Delete All Pending Values

Delete All pending Values button is enabled only when there are pending values based on the recent configuration changes. In case, user decides not to apply the configuration changes, user can click **Delete All Pending Values** button to terminate all the modifications. In case, the request fails to remove the BIOS attributes, it throws an error with corresponding HTTP Response Status code mapped to SMIL API error or Job Creation error. An EEMI message is generated and displayed at that point.

Pending Value

Configuration of a BIOS attribute via iDRAC is not applied immediately to BIOS. It requires a server reboot for the changes to take place. When you modify a BIOS attribute then **Pending Value** gets updated. If an attribute already has a pending value (and that has been configured) it is displayed on the GUI.

Modifying BIOS Configuration

Modifying BIOS configuration results in audit log entries, which gets entered in LC logs.

BIOS Live Scanning

BIOS live scanning verifies the integrity and authenticity of the BIOS image in the BIOS primary ROM when the host is powered ON but not in POST.

() NOTE:

- This feature requires iDRAC Datacenter license.
- You need to have Debug privilege for operating this feature.

iDRAC performs verification of immutable sections of BIOS image automatically at the followings scenarios:

- At AC cycle/Cold boot
- On a schedule determined by user
- On demand (initiated by user)

Successful result of live scanning is logged to LC log. Failure result is logged to both LCL and SEL.

Topics:

- BIOS Live Scanning
- BIOS Recovery and Hardware Root of Trust (RoT)

BIOS Live Scanning

BIOS live scanning verifies the integrity and authenticity of the BIOS image in the BIOS primary ROM when the host is powered ON but not in POST.

() NOTE:

- This feature requires iDRAC Datacenter license.
- You need to have Debug privilege for operating this feature.

iDRAC performs verification of immutable sections of BIOS image automatically at the followings scenarios:

- At AC cycle/Cold boot
- On a schedule determined by user
- On demand (initiated by user)

Successful result of live scanning is logged to LC log. Failure result is logged to both LCL and SEL.


BIOS Recovery and Hardware Root of Trust (RoT)

For PowerEdge server, it is mandatory to recover from corrupted or damaged BIOS image either due to malicious attack or power surges or any other unforeseeable events. An alternate reserve of BIOS image would be necessary to recover BIOS in order to bring the PowerEdge server back to functional mode from unbootable mode. This alternative/recovery BIOS is stored in a 2nd SPI (mux'ed with primary BIOS SPI).

The recovery sequence can be initiated through any of the following approaches with iDRAC as the main orchestrator of the BIOS recovery task:

- 1. Auto recovery of BIOS primary image/recovery image BIOS image is recovered automatically during the host boot process after the BIOS corruption is detected by BIOS itself.
- 2. Forced recovery of BIOS Primary/recovery image User initiates an OOB request to update BIOS either because they have a new updated BIOS or BIOS was just crashing by failing to boot.
- **3. Primary BIOS ROM update** The single Primary ROM is split into Data ROM and Code ROM. iDRAC has full access/ control over Code ROM. It switches MUX to access Code ROM whenever needed.
- 4. BIOS Hardware Root of Trust (RoT) This feature is available in severs with model number RX5X, CX5XX, and TX5X. During every host boot (only cold boot or A/C cycle, not during warm reboot), iDRAC ensures that RoT is performed. RoT runs automatically and user cannot initiate it using any interfaces. This iDRAC boot first policy verifies host BIOS ROM contents on every AC cycle and host DC cycle. This process ensures secure boot of BIOS and further secures the host boot process.

NOTE: For more information on Hardware RoT, refer to this link: https://downloads.dell.com/Manuals/Common/dell-emcidrac9-security-root-of-trust-bios-live-scanning.pdf



Configuring and using virtual console

iDRAC has added an enhanced HTML5 option in vConsole which allows vKVM (virtual Keyboard, Video, and Mouse) over standard VNC client. You can use the virtual console to manage a remote system using the keyboard, video, and mouse on your management station to control the corresponding devices on a managed server. This is a licensed feature for rack and tower servers. It is available by default in blade servers. You need iDRAC Configure privilege to access all configurations on virtual console.

Following are the list of configurable attributes in Virtual Console:

- vConsole Enabled Enabled / Disabled
- Max Sessions 1-6
- Active sessions 0-6
- Remote Presence Port (Not applicable for eHTML5 plugin)
- Video Encryption Enabled / Disabled (Not applicable for eHTML5 plugin)
- Local Server Video Enabled / Disabled
- Plug-in Type eHTML5 (by default), ActiveX, Java, HTML5
- Dynamic Action on Sharing Request Timeout Full Access, Read Only Access, And Deny Access
- Automatic System Lock Enabled / Disabled
- Keyboard/Mouse Attach State Auto-attach, Attached, and Detached

The key features are:

- A maximum of six simultaneous Virtual Console sessions are supported. All the sessions view the same managed server console simultaneously.
- You can launch virtual console in a supported web browser by using Java, ActiveX, HTML5, or eHTML5 plug-in.
 NOTE: By default, the virtual console type is set to eHTML5.

(i) NOTE: Any change in web server configuration will result in termination of existing virtual console session.

- When you open a Virtual Console session, the managed server does not indicate that the console has been redirected.
- You can open multiple Virtual Console sessions from a single management station to one or more managed systems simultaneously.
- You cannot open two virtual console sessions from the management station to the managed server using the same HTML5 plug-in.
- If a second user requests a Virtual Console session, the first user is notified and is given the option to refuse access, allow read-only access, or allow full shared access. The second user is notified that another user has control. The first user must respond within thirty seconds, or else access is granted to the second user based on the default setting. If neither the first or second user has administrator privileges, terminating the first user's session automatically terminates the second user's session.
- Boot logs and crash logs are captured as Video logs and are in MPEG1 format.
- Crash screen is captured as JPEG file.
- Keyboard macros are supported on all plug-ins.
- Keyboard macros are supported on all plug-ins. Following are the list of macros that are supported by ActiveX and Java plug-ins:

Table 57. Keyboard Macros Supported by ActiveX and Java plug-ins

MAC Client	Win Client	Linux Client
Ctrl-Al-Del	Ctrl-Alt-Del	Ctrl-Alt-Del
Alt-SysRq-B	Alt-SysRq-B	Alt-SysRq-B
-	Win-P	-
-	-	Ctrl-Alt-F<1-12>
Alt-SysRq	-	-



Table 57. Keyboard Macros Supported by ActiveX and Java plug-ins (continued)

MAC Client	Win Client	Linux Client
SysRq	-	-
PrtScrn	-	-
Alt-PrtScrn	-	-
Pause	-	-

(i) NOTE: For keyboard macros supported in HTML plug-in, see the section HTML5 based virtual console.

NOTE: The number of active virtual-console sessions displayed in the web interface is only for active web-interface sessions. This number does not include sessions from other interfaces such as SSH and RACADM.

NOTE: For information about configuring your browser to access the virtual console, see Configuring web browsers to use virtual console.

(i) NOTE: To disable KVM access, use the **Disable** option under the settings for chassis in the OME Modular web interface.

Topics:

- Supported screen resolutions and refresh rates
- Configuring virtual console
- Previewing virtual console
- Launching virtual console
- Using virtual console viewer

Supported screen resolutions and refresh rates

The following table lists the supported screen resolutions and corresponding refresh rates for a Virtual Console session running on the managed server.

Table 58. Supported screen resolutions and refresh rates

Screen Resolution	Refresh Rate (Hz)
720x400	70
640x480	60, 72, 75, 85
800×600	60, 70, 72, 75, 85
1024x768	60, 70, 72, 75, 85
1280x1024	60
1920x1200	60

It is recommended that you configure the monitor display resolution to 1920x1200 pixels.

Virtual Console supports a maximum video resolution of 1920x1200 at 60 Hz refresh rate. In order to achieve this resolution, following conditions are required:

- KVM / monitor attached to VGA that supports 1920x1200 resolution
- Latest Matrox video driver (for Windows)

When a local KVM / Monitor with maximum resolution below 1920x1200 is connected to either VGA connector, it will reduce the maximum resolution supported in virtual console.

iDRAC virtual console leverages the onboard Matrox G200 graphics controller to determine the maximum resolution of the attached monitor when a physical display is present. When the monitor supports 1920x1200 or greater resolution, the virtual



console supports 1920x1200 resolution. If the monitor attached supports lower max resolution (like many KVMs), the virtual console max resolution is limited.

Maximum virtual console resolutions based on monitor display ratio:

- 16:10 monitor: 1920x1200 will be the max resolution
- 16:9 monitor: 1920x1080 will be the max resolution

When a physical monitor is not connected to either VGA port on the server, the OS installed will dictate the available resolutions for virtual console.

Maximum virtual console resolutions based on host OS without physical monitor:

- Windows: 1600x1200 (1600x1200, 1280x1024, 1152x864, 1024x768, 800x600)
- Linux: 1024x768 (1024x768, 800x600, 848x480, 640x480)

NOTE: If a higher resolution through virtual console is required when physical KVM or monitor is not present, a VGA Display Emulator dongle can be leveraged to mimic an external monitor connection with a resolution up to 1920x1080.

() NOTE: If you have an active Virtual Console session and a lower resolution monitor is connected to the Virtual Console, the server console resolution may reset if the server is selected on the local console. If the system is running a Linux operating system, an X11 console may not be viewable on the local monitor. Press <Ctrl><Alt><F1> at the iDRAC Virtual Console to switch Linux to a text console.

Configuring virtual console

Before configuring the Virtual Console, make sure that the management station is configured.

You can configure the virtual console using iDRAC Web interface or RACADM command line interface.

Configuring virtual console using web interface

To configure Virtual Console using iDRAC Web interface:

- 1. Go to Configuration > Virtual Console. Click Start the Virtual Console link, then Virtual Console page is displayed.
- 2. Enable virtual console and specify the required values. For information about the options, see the *iDRAC Online Help*.
 (i) NOTE: If you are using Nano operating system, disable the Automatic System Lock feature on the Virtual Console
- **3.** Click **Apply**. The virtual console is configured.

page.

Configuring virtual console using RACADM

To configure the Virtual Console, use the set command with the objects in the **iDRAC.VirtualConsole** group.

For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Previewing virtual console

Before launching the Virtual Console, you can preview the state of the Virtual Console on the **System** > **Properties** > **System Summary** page. The **Virtual Console Preview** section displays an image showing the state of the Virtual Console. The image is refreshed every 30 seconds. This is a licensed feature.

(i) NOTE: The Virtual Console image is available only if you have enabled Virtual Console.

Launching virtual console

You can launch the virtual console using the iDRAC Web Interface or a URL.

(i) NOTE: Do not launch a Virtual Console session from a Web browser on the managed system.



Before launching the Virtual Console, make sure that:

- You have administrator privileges.
- Web browser is configured to use HTML5, eHTML5, Java, or ActiveX plug-ins.
- Minimum network bandwidth of 1 MB/sec is available.

NOTE: If the embedded video controller is disabled in BIOS and if you launch the Virtual Console, the Virtual Console Viewer is blank.

While launching Virtual Console using 32-bit or 64-bit IE browsers, use HTML5/eHTML5, or use the required plug-in (Java or ActiveX) that is available in the respective browser. The Internet Options settings are common for all browsers.

While launching the Virtual Console using Java plug-in, occasionally you may see a Java compilation error. To resolve this, go to **Java control panel** > **General** > **Network Settings** and select **Direct Connection**.

If the Virtual Console is configured to use ActiveX plug-in, it may not launch the first time. This is because of the slow network connection and the temporary credentials (that Virtual Console uses to connect) timeout is two minutes. The ActiveX client plug-in download time may exceed this time. After the plug-in is successfully downloaded, you can launch the Virtual Console normally.

To launch the Virtual Console by using HTML5/eHTML5 plug-in, you must disable the pop-up blocker.

Virtual Console has the following Console controls:

- 1. General You can set Keyboard Macros, Aspect Ratio, and Touch Mode.
- 2. ${\rm KVM}$ Shows the values for Frame Rate, Bandwidth, Compression, and Packet Rate.
- **3. Performance** You can change the Video quality and Video speed using this option.
- 4. User List You can view the list of users connected to the console.

You can access Virtual Media by clicking the Connect to Virtual Media option available in virtual console.

Launching virtual console using web interface

You can launch the virtual console in the following ways:

• Go to Configuration > Virtual Console. Click Start the Virtual Console link. Virtual console page is displayed.

The **Virtual Console Viewer** displays the remote system's desktop. Using this viewer, you can control the remote system's mouse and keyboard functions from your management station.

Multiple message boxes may appear after you launch the application. To prevent unauthorized access to the application, navigate through these message boxes within three minutes. Otherwise, you are prompted to relaunch the application.

If one or more Security Alert windows appear while launching the viewer, click Yes to continue.

Two mouse pointers may appear in the viewer window: one for the managed server and another for your management station.

Launching virtual console using a URL

To launch the Virtual Console using the URL:

- 1. Open a supported Web browser and in the address box, type the following URL in lower case: https://iDRAC_ip/console
- 2. Based on the login configuration, the corresponding **Login** page is displayed:
 - If Single Sign On is disabled and Local, Active Directory, LDAP, or Smart Card login is enabled, the corresponding **Login** page is displayed.
 - If Single-Sign On is enabled, the Virtual Console Viewer is launched and the Virtual Console page is displayed in the background.
 - () NOTE: Internet Explorer supports Local, Active Directory, LDAP, Smart Card (SC) and Single Sign-On (SSO) logins. Firefox supports Local, AD, and SSO logins on Windows-based operating system and Local, Active Directory, and LDAP logins on Linux-based operating systems.
 - **NOTE:** If you do not have Access Virtual Console privilege but have Access Virtual Media privilege, then using this URL launches the Virtual Media instead of the Virtual Console.



Disabling warning messages while launching virtual console or virtual media using Java or ActiveX plug-in

You can disable the warning messages while launching the Virtual Console or Virtual Media using Java plug-in.

(i) NOTE: You need Java 8 or later to use this feature and to launch iDRAC Virtual Console over an IPv6 network.

 Initially, when you launch Virtual Console or Virtual Media using Java plug-in, the prompt to verify the publisher is displayed. Click Yes.

A certificate warning message is displayed indicating that a trusted certificate is not found.

NOTE: If the certificate is found in the operating system's certificate store or if it is found in a previously specified user location, then this warning message is not displayed.

2. Click Continue.

The Virtual Console Viewer or Virtual Media Viewer is launched.

(i) NOTE: The Virtual Media viewer is launched if Virtual Console is disabled.

- 3. From the **Tools** menu, click **Session Options** and then **Certificate** tab.
- 4. Click Browse Path, specify the location to store the user's certificate, click Apply, click OK, and exit from the viewer.
- 5. Launch Virtual Console again.
- 6. In the certificate warning message, select the Always trust this certificate option, and then click Continue.
- 7. Exit from the viewer.
- 8. When you re-launch Virtual Console, the warning message is not displayed.

Using virtual console viewer

The Virtual Console Viewer provides various controls such as mouse synchronization, virtual console scaling, chat options, keyboard macros, power actions, next boot devices, and access to Virtual Media. For information to use these features, see the *iDRAC Online Help*.

(i) NOTE: If the remote server is powered off, the message 'No Signal' is displayed.

The Virtual Console Viewer title bar displays the DNS name or the IP address of the iDRAC you are connected to from the management station. If iDRAC does not have a DNS name, then the IP address is displayed. The format is:

• For rack and tower servers:

<DNS name / IPv6 address / IPv4 address>, <Model>, User: <username>, <fps>

• For blade servers:

<DNS name / IPv6 address / IPv4 address>, <Model>, <Slot number>, User: <username>, <fps>

Sometimes the Virtual Console Viewer may display low quality video. This is due to slow network connectivity that leads to loss of one or two video frames when you start the Virtual Console session. To transmit all the video frames and improve the subsequent video quality, do any of the following:

- In the System Summary page, under Virtual Console Preview section, click Refresh.
- In the Virtual Console Viewer, under Performance tab, set the slider to Maximum Video Quality.

eHTML5 based virtual console

NOTE: While using eHTML5 to access virtual console, the language must be consistent across client and target keyboard layout, OS, and browser. For example, all must be in English (US) or any of the supported languages.

To launch the eHTML5 virtual console, you must enable the virtual console feature from the iDRAC Virtual Console page and set the **Plug-in Type** option to eHTML5.

(i) NOTE: By default the virtual console type is set to eHTML5.

You can launch virtual console as a pop-up window by using one of the following methods:

• From iDRAC Home page, click the **Start the Virtual Console** link available in the Console Preview session

278 Configuring and using virtual console



- From iDRAC Virtual Console page, click Start the Virtual Console link.
- From iDRAC login page, type https//<iDRAC IP>/console. This method is called as Direct Launch.

In the eHTML5 virtual console, the following menu options are available:

- Power
- Boot
- Chat
- Keyboard
- Screen Capture
- Refresh
- Full Screen
- Disconnect Viewer
- Console Controls
- Virtual Media

The **Pass all keystrokes to server** option is not supported on eHTML5 virtual console. Use keyboard and keyboard macros for all the functional keys.

- General
 - **Console control** This has the following configuration options:
 - Keyboard Macros This is supported in eHTML5 virtual console and are listed as the following drop-down options. Click **Apply** to apply the selected key combination on the server.
 - Ctrl+Alt+Del
 - Ctrl+Alt+F1
 - Ctrl+Alt+F2
 - Ctrl+Alt+F3
 - Ctrl+Alt+F4
 - Ctrl+Alt+F5
 - Ctrl+Alt+F6
 - Ctrl+Alt+F7
 - Ctrl+Alt+F8
 - Ctrl+Alt+F9
 - Ctrl+Alt+F10
 - Ctrl+Alt+F11
 - Ctrl+Alt+F12
 - Alt+Tab
 - Alt+ESC
 - Ctrl+ESC
 - Alt+Space
 - Alt+Enter
 - Alt+Hyphen
 - Alt+F1
 - Alt+F2
 - Alt+F3
 - Alt+F4
 - Alt+F5
 - Alt+F6
 - Alt+F7
 - Alt+F8
 - Alt+F9
 - Alt+F10
 - Alt+F11
 - Alt+F12
 - PrntScrn
 - Alt+PrntScrn
 - F1
 - Pause
 - Tab
 - Ctrl+Enter



- SysRq
- Alt+SysRq
- Win-P
- Aspect Ratio The eHTML5 virtual console video image automatically adjusts the size to make the image visible. The following configuration options are displayed as a drop-down list:
 - Maintain
 - Don't Maintain

Click **Apply** to apply the selected settings on the server.

- Touch Mode The eHTML5 virtual console supports the Touch Mode feature. The following configuration options are displayed as a drop-down list:
 - Direct
 - Relative

Click **Apply** to apply the selected settings on the server.

- **Virtual Clipboard** Virtual clipboard enables you to cut / copy / paste text buffer from virtual console to iDRAC host server. Host server could be BIOS, UEFI or in OS prompt. This is a one-way action from client computer to iDRAC's host server only. Follow these steps to use the Virtual clipboard:
 - Place the mouse cursor or keyboard focus on the desired window in the host server desktop.
 - Select the **Console Controls** menu from vConsole.
 - Copy the OS clipboard buffer using keyboard hotkeys, mouse, or touch pad controls depending on the Client OS. Or, you can type the text manually in the text box.

• Click Send Clipboard to Host.

 \circ $\;$ Then, the text appears on the host server's active window.

(i) NOTE:

- This feature is only available in Datacenter license.
- This feature only supports ASCII text.
- Control characters are not supported.
- Characters such as **New line** and **Tab** are allowed.
- Text buffer size is limited to 4000 characters.
- If more than maximum buffer is pasted, then the edit box in iDRAC GUI will truncate it to maximum buffer size.
- **KVM** This menu has list of the following read only components:
- Frame Rate
- Bandwidth
- Compression
- Packet Rate
- Performance You can use the slider button to adjust Maximum Video Quality and Maximum Video Speed.
- User List You can see the list of users that are logged in to the Virtual console.
- **Keyboard** The difference between physical and virtual keyboard is that virtual keyboard changes its layout according to the browser language.
- Virtual Media Click Connect Virtual Media option to start the virtual media session.
- **Connect Virtual Media** This menu contains the options for Map CD/DVD, Map Removable Disk, Map External Device, and Reset USB.
 - Virtual Media Statistics This menu shows the Transfer Rate (Read-only). Also, it shows the details of CD/DVD and Removable Disks details such as Mapping details, status (read-only or not), duration, and Read/Write Bytes.
- **Create Image** This menu allows you to select a local folder and generate FolderName.img file with local folder contents.

NOTE: For security reasons read/write access is disabled while accessing virtual console in eHTML5. With Java or ActiveX plug-ins, you can accept security messaging before the plug-in is given the read/write authority.

Supported Browsers

The eHTML5 virtual console is supported on the following browsers:

- Internet Explorer 11
- Chrome 78/79
- Firefox 70/71



• Safari 13.1

(i) NOTE: It is recommended to have Mac OS version 10.10.2 (or onward) installed in the system.

For more details on supported browsers and versions, see the *iDRAC Release Notes* available at https://www.dell.com/ idracmanuals.

HTML5 based virtual console

() NOTE: While using HTML5 to access virtual console, the language must be consistent across client and target keyboard layout, OS, and browser. For example, all must be in English (US) or any of the supported languages.

To launch the HTML5 virtual console, you must enable the virtual console feature from the iDRAC Virtual Console page and set the **Plug-in Type** option to HTML5.

You can launch virtual console as a pop-up window by using one of the following methods:

- From iDRAC Home page, click the Start the Virtual Console link available in the Console Preview session
- From iDRAC Virtual Console page, click Start the Virtual Console link.
- From iDRAC login page, type https//<iDRAC IP>/console. This method is called as Direct Launch.

In the HTML5 virtual console, the following menu options are available:

- Power
- Boot
- Chat
- Keyboard
- Screen Capture
- Refresh
- Full Screen
- Disconnect Viewer
- Console Controls
- Virtual Media

The **Pass all keystrokes to server** option is not supported on HTML5 virtual console. Use keyboard and keyboard macros for all the functional keys.

- **Console control** This has the following configuration options:
 - Keyboard Macros This is supported in HTML5 virtual console and are listed as the following drop-down options. Click **Apply** to apply the selected key combination on the server.
 - Ctrl+Alt+Del
 - Ctrl+Alt+F1
 - Ctrl+Alt+F2
 - Ctrl+Alt+F3
 - Ctrl+Alt+F4
 - Ctrl+Alt+F5
 - Ctrl+Alt+F6
 - Ctrl+Alt+F7
 - Ctrl+Alt+F8
 - Ctrl+Alt+F9
 - Ctrl+Alt+F10
 - Ctrl+Alt+F11
 - Ctrl+Alt+F12
 - Alt+Tab
 - Alt+ESC
 - Ctrl+ESC
 - Alt+Space
 - Alt+Enter
 - Alt+Hyphen
 - Alt+F1
 - Alt+F2
 - Alt+F3

Configuring and using virtual console 281



- Alt+F4
- Alt+F5
- Alt+F6
- Alt+F7
- Alt+F8
- Alt+F9
- Alt+F10
- Alt+F11
- Alt+F12
- PrntScrn
- Alt+PrntScrn
- F1
- Pause
- Tab
- Ctrl+Enter
- SysRq
- Alt+SysRq
- Win-P
- Aspect Ratio The HTML5 virtual console video image automatically adjusts the size to make the image visible. The following configuration options are displayed as a drop-down list:
 - Maintain
 - Don't Maintain

Click **Apply** to apply the selected settings on the server.

- Touch Mode The HTML5 virtual console supports the Touch Mode feature. The following configuration options are displayed as a drop-down list:
 - Direct
 - Relative

Click **Apply** to apply the selected settings on the server.

- **Virtual Clipboard** Virtual clipboard enables you to cut / copy / paste text buffer from virtual console to iDRAC host server. Host server could be BIOS, UEFI or in OS prompt. This is a one-way action from client computer to iDRAC's host server only. Follow these steps to use the Virtual clipboard:
 - Place the mouse cursor or keyboard focus on the desired window in the host server desktop.
 - Select the **Console Controls** menu from vConsole.
 - Copy the OS clipboard buffer using keyboard hotkeys, mouse, or touch pad controls depending on the Client OS. Or, you can type the text manually in the text box.
 - Click Send Clipboard to Host.
 - Then, the text appears on the host server's active window.

() NOTE:

- This feature is only available in Datacenter license.
- This feature only supports ASCII text.
- Control characters are not supported.
- Characters such as **New line** and **Tab** are allowed.
- Text buffer size is limited to 4000 characters.
- If more than maximum buffer is pasted, then the edit box in iDRAC GUI will truncate it to maximum buffer size.
- **Keyboard** The difference between physical and virtual keyboard is that virtual keyboard changes its layout according to the browser language.
- Touch Mode The HTML5 virtual console supports the Touch Mode feature. The following configuration options are displayed as a drop-down list:
 - o Direct

Relative

- Click **Apply** to apply the selected settings on the server.
- Mouse Acceleration Select the mouse acceleration based on the operating system. The following configuration options are displayed as a drop-down list:
 - Absolute (Windows, latest versions of Linux, Mac OS-X)



- Relative, no acceleration
- Relative (RHEL, earlier versions of Linux)
- Linux RHEL 6.x and SUSE Linux Enterprise Server 11 or later

Click **Apply** to apply the selected settings on the server.

• Virtual Media — Click Connect Virtual Media option to start the virtual media session. when the virtual media is connected, you can see the options like Map CD/DVD, Map Removable Disk, and Reset USB.

NOTE: For security reasons read/write access is disabled while accessing virtual console in HTML5. With Java or ActiveX plug-ins, you can accept security messaging before the plug-in is given the read/write authority.

Supported Browsers

The HTML5 virtual console is supported on the following browsers:

- Internet Explorer 11
- Chrome 78/79
- Firefox 70/71
- Safari 13.1

(i) NOTE: It is recommended to have Mac OS version 10.10.2 (or onward) installed in the system.

For more details on supported browsers and versions, see the *iDRAC Release Notes* available at https://www.dell.com/ idracmanuals.

Synchronizing mouse pointers

(i) NOTE: This feature is not applicable with eHTML5 plugin type.

When you connect to a managed system through the Virtual Console, the mouse acceleration speed on the managed system may not synchronize with the mouse pointer on the management station and displays two mouse pointers in the Viewer window.

When using Red Hat Enterprise Linux or Novell SUSE Linux, configure the mouse mode for Linux before you launch the Virtual Console viewer. The operating system's default mouse settings are used to control the mouse arrow in the Virtual Console viewer.

When two mouse cursors are seen on the client Virtual Console viewer, it indicates that the server's operating system supports Relative Positioning. This is typical for Linux operating systems or Lifecycle Controller and causes two mouse cursors if the server's mouse acceleration settings are different from the mouse acceleration settings on the Virtual Console client. To resolve this, switch to single cursor or match the mouse acceleration on the managed system and the management station:

- To switch to single cursor, from the **Tools** menu, select **Single Cursor**.
- To set the mouse acceleration, go to Tools > Session Options > Mouse. Under Mouse Acceleration tab, select Windows
 or Linux based on the operating system.

To exit single cursor mode, press <F9> or the configured termination key.

NOTE: This is not applicable for managed systems running Windows operating system since they support Absolute Positioning.

When using the Virtual Console to connect to a managed system with a recent Linux distribution operating system installed, you may experience mouse synchronization problems. This may be due to the Predictable Pointer Acceleration feature of the GNOME desktop. For correct mouse synchronization in the iDRAC Virtual Console, this feature must be disabled. To disable Predictable Pointer Acceleration, in the mouse section of the **/etc/X11/xorg.conf** file, add:

Option "AccelerationScheme" "lightweight".

If synchronization problems continue, do the following additional change in the <user_home>/.gconf/desktop/gnome/ peripherals/mouse/%gconf.xml file:

Change the values for motion_threshold and motion_acceleration to –1.

If you turn off mouse acceleration in GNOME desktop, in the Virtual Console viewer, go to **Tools** > **Session Options** > **Mouse**. Under **Mouse Acceleration** tab, select **None**.

For exclusive access to the managed server console, you must disable the local console and re-configure the **Max Sessions** to 1 on the **Virtual Console page**.

Configuring and using virtual console 283



Passing all keystrokes through virtual console for Java or ActiveX plug-in

You can enable the **Pass all keystrokes to server** option and send all keystrokes and key combinations from the management station to the managed system through the Virtual Console Viewer. If it is disabled, it directs all the key combinations to the management station where the Virtual Console session is running. To pass all keystrokes to the server, in the Virtual Console Viewer, go to **Tools** > **Session Options** > **General** tab and select the **Pass all keystrokes to server** option to pass the management station's keystrokes to the managed system.

The behavior of the Pass all keystrokes to server feature depends on the:

Plug-in type (Java or ActiveX) based on which Virtual Console session is launched.

For the Java client, the native library must be loaded for Pass all keystrokes to server and Single Cursor mode to function. If the native libraries are not loaded, the **Pass all keystrokes to server** and **Single Cursor** options are deselected. If you attempt to select either of these options, an error message is displayed indicating that the selected options are not supported.

For the ActiveX client, the native library must be loaded for Pass all keystrokes to server function to work. If the native libraries are not loaded, the **Pass all keystrokes to server** option is deselected. If you attempt to select this option, an error message is displayed indicating that the feature is not supported

For MAC operating systems, enable the **Enable access of assistive device** option in **Universal Access** for the Pass all keystrokes to server feature to work.

- Operating system running on the management station and managed system. The key combinations that are meaningful to the operating system on the management station are not passed to the managed system.
- Virtual Console Viewer mode—Windowed or Full Screen.

In Full Screen mode, **Pass all keystrokes to server** is enabled by default.

In Windowed mode, the keys passed only when the Virtual Console Viewer is visible and is active.

When changed from Full Screen mode to Windowed mode, the previous state of Pass all keys is resumed.

Java-based virtual console session running on Windows operating system

- Ctrl+Alt+Del key is not sent to the managed system, but always interpreted by the management station.
 - When Pass All Keystrokes to Server is enabled, the following keys are not sent to the managed system:
 - Browser Back Key
 - Browser Forward Key
 - Browser Refresh key
 - o Browser Stop Key
 - Browser Search Key
 - Browser Favorites key
 - Browser Start and Home key
 - Volume mute key
 - Volume down key
 - Volume up key
 - Next track key
 - Previous track key
 - Stop Media key
 - Play/Pause media key
 - o Start mail key
 - Select media key
 - Start Application 1 key
 - Start Application 2 key
- All the individual keys (not a combination of different keys, but a single key stroke) are always sent to the managed system. This includes all the Function keys, Shift, Alt, Ctrl key and Menu keys. Some of these keys affect both management station and managed system.

For example, if the management station and the managed system is running Windows operating system, and Pass All Keys is disabled, when you press the Windows key to open the **Start** Menu, the **Start** menu opens on both management station and



managed system. However, if Pass All Keys is enabled, then the **Start** menu is opened only on the managed system and not on the management station.

• When Pass All Keys is disabled, the behavior depends on the key combinations pressed and the special combinations interpreted by the operating system on the management station.

Java based virtual console session running on Linux operating system

The behavior mentioned for Windows operating system is also applicable for Linux operating system with the following exceptions:

- When Pass all keystrokes to server is enabled, <Ctrl+Alt+Del> is passed to the operating system on the managed system.
- Magic SysRq keys are key combinations interpreted by the Linux Kernel. It is useful if the operating system on the
 management station or the managed system freezes and you need to recover the system. You can enable the magic SysRq
 keys on the Linux operating system using one of the following methods:
 - Add an entry to **/etc/sysctl.conf**
 - echo "1" > /proc/sys/kernel/sysrq
- When Pass all keystrokes to server is enabled, the magic SysRq keys are sent to the operating system on the managed system. The key sequence behavior to reset the operating system, that is reboot without un-mounting or sync, depends on whether the magic SysRq is enabled or disabled on the management station:
 - If SysRq is enabled on the management station, then <Ctrl+Alt+SysRq+b> or <Alt+SysRq+b> resets the management station irrespective of the system's state.
 - If SysRq is disabled on the management station, then the <Ctrl+Alt+SysRq+b> or <Alt+SysRq+b>keys resets the operating system on the managed system.
 - Other SysRq key combinations (example, <Alt+SysRq+k>, <Ctrl+Alt+SysRq+m>, and so on) are passed to the managed system irrespective of the SysRq keys enabled or not on the management station.

Using SysRq magic keys through remote console

You can enable SysRq magic keys through the remote console using any of the following:

- Opensoure IPMI tool
- Using SSH or External Serial Connector

Using opensource IPMI tool

Make sure that BIOS/iDRAC settings supports console redirection using SOL.

1. At the command prompt, run the SOL activate command:

Ipmitool -I lanplus -H <ipaddr> -U <username> -P <passwd> sol activate

The SOL session is activated.

- 2. After the server boots to the operating system, the localhost.localdomain login prompt appears. Log in using the operating system user name and password.
- **3.** If SysRq is not enabled, enable using echo 1 >/proc/sys/kernel/sysrq.
- 4. Run break sequence ~B.
- **5.** Use the SysRq magic key to enable the SysRq function. For example, the following command displays the memory information on the console:

echo m > /proc/sysrq-trigger displays

Using SSH or external serial connector directly connecting through serial cable

- 1. For SSH sessions, after logging in using the iDRAC username and password, at the /admin> prompt, run the command console com2. The localhost.localdomain prompt appears.
- 2. For console redirection using external serial connector directly connected to the system through a serial cable, the localhost.localdomain login prompt appears after the server boots to the operating system.
- 3. Log in using the operating system user name and password.
- 4. If SysRq is not enabled, enable using echo 1 >/proc/sys/kernel/sysrq.



5. Use the magic key to enable the SysRg function. For example, the following command reboots the server:

echo b > /proc/sysrq-trigger

(i) NOTE: You do not have to run break sequence before using the magic SysRq keys.

ActiveX based virtual console session running on Windows operating system

The behavior of the pass all keystrokes to server feature in ActiveX based Virtual Console session running on Windows operating system is similar to the behavior explained for Java based Virtual Console session running on the Windows management station with the following exceptions:

• When Pass All Keys is disabled, pressing F1 launches the application Help on both management station and managed system, and the following message is displayed:

Click Help on the Virtual Console page to view the online Help

- The media keys may not be blocked explicitly.
- <Alt + Space>, <Ctrl + Alt + +>, <Ctrl + Alt + -> are not sent to the managed system and is interpreted by the operating system on the management station.



Using iDRAC Service Module

The iDRAC Service Module is a software application that is recommended to be installed on the server (it is not installed by default). It complements iDRAC with monitoring information from the operating system. It complements iDRAC by providing additional data to work with iDRAC interfaces such as the Web interface, Redfish, RACADM, and WSMan. You can configure the features monitored by the iDRAC Service Module to control the CPU and memory consumed on the server's operating system. Host OS command line interface has been introduced to enable or disable status of Full Power Cycle for all System components except the PSU.

(i) NOTE: iDRAC9 uses iSM version 3.01 and higher.

NOTE: You can use the iDRAC Service Module only if you have installed iDRAC Express or iDRAC Enterprise/Datacenter license.

Before using iDRAC Service Module, ensure that:

- You have login, configure, and server control privileges in iDRAC to enable or disable the iDRAC Service Module features.
- You do not disable the **iDRAC Configuration using local RACADM** option.
- OS to iDRAC pass-through channel is enabled through the internal USB bus in iDRAC.

(i) NOTE: If you perform LC wipe, idrac.Servicemodule values may still show the old values.

() NOTE:

- When iDRAC Service Module runs for the first time, by default it enables the OS to iDRAC pass-through channel in iDRAC. If you disable this feature after installing the iDRAC Service Module, then you must enable it manually in iDRAC.
- If the OS to iDRAC pass-through channel is enabled through LOM in iDRAC, then you cannot use the iDRAC Service Module.

Topics:

- Installing iDRAC Service Module
- Supported operating systems for iDRAC Service Module
- iDRAC Service Module monitoring features
- Using iDRAC Service Module from iDRAC web interface
- Using iDRAC Service Module from RACADM

Installing iDRAC Service Module

You can download and install the iDRAC Service Module from **dell.com/support**. You must have administrator privilege on the server's operating system to install the iDRAC Service Module. For information on installation, see the iDRAC Service Module User's Guide available at www.dell.com/idracservicemodule.

(i) NOTE: This feature is not applicable for Dell Precision PR7910 systems.

Installing iDRAC Service Module from iDRAC Express and Basic

From the iDRAC Service Module Setup page, click Install Service Module.

- 1. The Service Module Installer is available to the host operating system and a job is created in iDRAC. For Microsoft Windows operating system or Linux operating system, log in to the server either remotely or locally.
- 2. Find the mounted volume labeled as "SMINST" on your device list and run the appropriate script:
 - On Windows, open the command prompt and run the **ISM-Win.bat** batch file.
 - On Linux, open the shell prompt and run the ISM-Lx.sh script file.
- 3. After the installation is complete, iDRAC displays the Service Module as Installed and the installation date.



NOTE: The installer will be available to the host operating system for 30 minutes. If you do not start the installation within 30 minutes, you must restart the Service Module installation.

Installing iDRAC Service Module from iDRAC Enterprise

- 1. On the SupportAssist Registration wizard, click Next.
- 2. On the iDRAC Service Module Setup page, click Install Service Module.
- 3. Click Launch Virtual Console and click Continue on the security warning dialog box.
- 4. To locate the iSM installer file, log in to the server either remotely or locally.

(i) **NOTE:** The installer will be available to the host operating system for 30 minutes. If you do not start the installation within 30 minutes, you must restart the installation.

- 5. Find the mounted volume labeled as "SMINST" on your device list and run the appropriate script:
 - On Windows, open the command prompt and run the **ISM-Win.bat** batch file.
 - On Linux, open the shell prompt and run the **ISM-Lx.sh** script file.
- Follow the instructions on the screen to complete the installation. On the iDRAC Service Module Setup page, the Install Service Module button is disabled after the installation is complete and the Service Module status is displayed as Running.

Supported operating systems for iDRAC Service Module

For the list of operating systems supported by the iDRAC Service Module, see the iDRAC Service Module User's Guide available at www.dell.com/idracservicemodule.

iDRAC Service Module monitoring features

The iDRAC Service Module (iSM) provides the following monitoring features:

- Redfish profile support for network attributes
- iDRAC Hard Reset
- iDRAC access via Host OS (Experimental Feature)
- In-band iDRAC SNMP alerts
- View operating system (OS) information
- Replicate Lifecycle Controller logs to operating system logs
- Perform automatic system recovery options
- Populate Windows Management Instrumentation (WMI) Management Providers
- Integrate with SupportAssist Collection. This is applicable only if iDRAC Service Module version 2.0 or later is installed.
- Prepare to Remove NVMe PCIe SSD. For more informationhttps://www.dell.com/support/article/sln310557 .
- Remote Server Power Cycle

Redfish profile support for network attributes

iDRAC Service Module v2.3 or later provides additional network attributes to iDRAC, which can be obtained through the REST clients from iDRAC. For more details, see iDRAC Redfish profile support.

Operating system information

The OpenManage Server Administrator currently shares operating system information and host name with iDRAC. The iDRAC Service Module provides similar information such as OS name, OS version, and Fully Qualified Domain Name (FQDN) with



iDRAC. By default, this monitoring feature is enabled. It is not disabled if OpenManage Server Administrator is installed on the host OS.

In iSM version 2.0 or later, the operating system information feature is amended with the OS network interface monitoring. When iDRAC Service Module version 2.0 or later is used with iDRAC 2.00.00.00, it starts monitoring the operating system network interfaces. You can view this information using iDRAC web interface, RACADM, or WSMan.

Replicate Lifecycle logs to OS log

You can replicate the Lifecycle Controller Logs to the OS logs from the time when the feature is enabled in iDRAC. This is similar to the System Event Log (SEL) replication performed by OpenManage Server Administrator. All events that have the **OS Log** option selected as the target (in the **Alerts** page, or in the equivalent RACADM or WSMan interfaces) are replicated in the OS log using the iDRAC Service Module. The default set of logs to be included in the OS logs is the same as configured for SNMP alerts or traps.

iDRAC Service Module also logs the events that have occurred when the operating system is not functioning. The OS logging performed by iDRAC Service Module follows the IETF syslog standards for Linux-based operating systems.

() NOTE: Starting iDRAC Service Module version 2.1, the Lifecycle Controller Logs replication location in the Windows OS logs can be configured using the iDRAC Service Module installer. You can configure the location while installing iDRAC Service Module or modifying the iDRAC Service Module installer.

If OpenManage Server Administrator is installed, this monitoring feature is disabled to avoid duplicate SEL entries in the OS log.

NOTE: On Microsoft Windows, if iSM events get logged under System logs instead of Application logs, restart the Windows Event Log service or restart the host OS.

Automatic system recovery options

The Automatic system recovery feature is a hardware-based timer. If a hardware failure occurs, a notification may not be available, but the server is reset as if the power switch was activated. ASR is implemented using a timer that continuously counts down. The Health Monitor frequently reloads the counter to prevent it from counting down to zero. If the ASR counts down to zero, it is assumed that the operating system has locked up and the system automatically attempts to reboot.

You can perform automatic system recovery operations such as reboot, power cycle, or power off the server after a specified time interval. This feature is enabled only if the operating system watchdog timer is disabled. If OpenManage Server Administrator is installed, this monitoring feature is disabled to avoid duplicate watchdog timers.

Windows Management Instrumentation providers

WMI is a set of extensions to the Windows Driver Model that provides an operating system interface through which instrumented components provide information and notification. WMI is Microsoft's implementation of the Web-Based Enterprise Management (WBEM) and Common Information Model (CIM) standards from the Distributed Management Task Force (DMTF) to manage Server hardware, operating systems and applications. WMI Providers helps to integrate with Systems Management Consoles such as Microsoft System Center and enables scripting to manage Microsoft Windows Servers.

You can enable or disable the WMI option in iDRAC. iDRAC exposes the WMI classes through the iDRAC Service Module providing the server's health information. By default, WMI information feature is enabled. The iDRAC Service Module exposes the WSMan monitored classes in iDRAC through WMI. The classes are exposed in the root/cimv2/dcim namespace.

The classes can be accessed using any of the standard WMI client interfaces. For more information, see the profile documents.

This content use the **DCIM_iDRACCardString** and **DCIM_iDRACCardInteger** classes to illustrate the capability that WMI information feature provides in iDRAC Service Module. For the details of the supported classes and profiles, see the WSMan profiles documentation available at https://www.dell.com/support.

The listed attributes are used to configure **User Accounts** along with the required privileges:

AttributeName	WSMAN-Class	Privilege	License	Description	Supported Operation
UserName	DCIM_iDRACCardS tring	Write Privileges: ConfigUsers, Login	Basic	16users:	Enum, Get, Invoke

AttributeName	WSMAN-Class	Privilege	License	Description	Supported Operation
		Read Privileges:Login		Users.1#UserName to Users.16#UserNam e	
Password	DCIM_iDRACCardS tring	Write Privileges: ConfigUsers, Login Read Privileges: Login	Basic	Users.1#Password to Users.16#Password	Enum, Get, Invoke
Privilege	DCIM_iDRACCardl nteger	Write Privileges: ConfigUsers, Login Read Privileges: Login	Basic	Users.1#Password to Users.16#Password	Enum, Get, Invoke

- Enumerate or Get operation on the mentioned classes will provide the attribute related data.
- The attribute can be set by invoking the ApplyAttribute or SetAttribute command from **DCIM_iDRACCardService** class.
- () NOTE: DCIM_Account class is removed from WSMAN and provided the feature through attribute model.
 - **DCIM_iDRACCardString** and **DCIM_iDRACCardInteger** classes provide similar support to configure iDRAC users accounts.

Remote iDRAC Hard Reset

By using iDRAC, you can monitor the supported servers for critical system hardware, firmware, or software issues. Sometimes, iDRAC may become unresponsive due to various reasons. During such scenarios, you must turn off the server and reset iDRAC. To reset the iDRAC CPU, you must either power off and power on the server or perform an AC power cycle.

By using the remote iDRAC hard reset feature, whenever iDRAC becomes unresponsive, you can perform a remote iDRAC reset operation without an AC power cycle.. To reset the iDRAC remotely, make sure that you have administrative privileges on the host OS. By default, the remote iDRAC hard reset feature is enabled. You can perform a remote iDRAC hard reset using iDRAC Web interface, RACADM, and WSMan.

Command usage

This section provides the command usages for Windows, Linux, and ESXi operating systems to perform iDRAC hard reset.

Windows

- Using the local Windows Management Instrumentation (WMI):
- o winrm i iDRACHardReset wmi/root/cimv2/dcim/DCIM_iSMService? InstanceID="iSMExportedFunctions"
- Using the remote WMI interface:

```
winrm i iDRACHardReset wmi/root/cimv2/dcim/dcim_ismservice?
InstanceID="iSMExportedFunctions" -u:<admin-username> -p:<admin-password> -r:http://
<remote-hostname OR IP>/wsman -a:Basic -encoding:utf-8 -skipCACheck -skipCNCheck
```

• Using the Windows PowerShell script with force and without force:

```
Invoke-iDRACHardReset -force
```

Invoke-iDRACHardReset

• Using the **Program Menu** shortcut:

For simplicity, iSM provides a shortcut in the **Program Menu** of the Windows operating system. When you select the **Remote iDRAC Hard Reset** option, you are prompted for a confirmation to reset the iDRAC. After you confirm, the iDRAC is reset and the result of the operation is displayed.

NOTE: The following warning message appears in the **Event Viewer** under the **Application Logs** category. This warning does not require any further action.



() NOTE: A provider, ismserviceprovider, has been registered in the Windows Management Instrumentation namespace Root\CIMV2\DCIM to use the LocalSystem account. This account is privileged and the provider may cause a security violation if it does not correctly impersonate user requests.

• Linux

iSM provides an executable command on all iSM supported Linux operating system. You can run this command by logging into the operating system by using SSH or equivalent.

Invoke-iDRACHardReset

Invoke-iDRACHardReset -f

ESXi

On all iSM supported ESXi operating systems, the iSM v2.3 supports a Common Management Programming Interface (CMPI) method provider to perform the iDRAC reset remotely by using the WinRM remote commands.

```
winrm i iDRACHardReset http://schemas.dell.com/wbem/wscim/1/cim-schema/2/root/cimv2/dcim/
DCIM_iSMService?__cimnamespace=root/cimv2/dcim+InstanceID= iSMExportedFunctions -u:<root-
username> -p:<passwd> -r:https://<Host-IP>:443/WSMan -a:basic -encoding:utf-8
-skipCNCheck -skipCACheck -skipRevocationcheck
```

(i) NOTE: VMware ESXi operating system does not prompt for confirmation before resetting the iDRAC.

NOTE: Due to limitations on the VMware ESXi operating system, iDRAC connectivity is not restored completely after the reset. Ensure that you manually reset iDRAC.

Table 59. Error Handling

Result	Description
0	Success
1	Unsupported BIOS version for iDRAC reset
2	Unsupported platform
3	Access denied
4	iDRAC reset failed

In-band Support for iDRAC SNMP Alerts

By using iDRAC Service Module v2.3, you can receive SNMP alerts from the host operating system, which is similar to the alerts that are generated by iDRAC.

You can also monitor the iDRAC SNMP alerts without configuring the iDRAC and manage the server remotely by configuring the SNMP traps and destination on the host OS. In iDRAC Service Module v2.3 or later, this feature converts all the Lifecycle logs replicated in the OS logs into SNMP traps.

(i) NOTE: This feature is active only if the Lifecycle Logs replication feature is enabled.

NOTE: On Linux operating systems, this feature requires a master or OS SNMP enabled with SNMP multiplexing (SMUX) protocol.

By default, this feature is disabled. Though the In-band SNMP alerting mechanism can coexist along with iDRAC SNMP alerting mechanism, the recorded logs may have redundant SNMP alerts from both the sources. It is recommended to either use the in-band or out-of-band option, instead of using both.

Command usage

This section provides the command usages for Windows, Linux, and ESXi operating systems.

• Windows operating system

• Using the local Windows Management Instrumentation (WMI):

```
winrm i EnableInBandSNMPTraps wmi/root/cimv2/dcim/DCIM_iSMService?
InstanceID="iSMExportedFunctions" @{state="[0/1]"}
```



• Using the remote WMI interface:

```
winrm i EnableInBandSNMPTraps wmi/root/cimv2/dcim/DCIM_iSMService?
InstanceID="iSMExportedFunctions" @{state="[0/1]"} -u:<admin-username> -p:<admin-
passwd> -r:http://<remote-hostname OR IP>/WSMan -a:Basic -encoding:utf-8 -skipCACheck
-skipCNCheck
```

Linux operating system

On all iSM supported Linux operating system, iSM provides an executable command. You can run this command by logging into the operating system by using SSH or equivalent.

Beginning with iSM 2.4.0, you can configure Agent-x as the default protocol for in-band iDRAC SNMP alerts using the following command:

./Enable-iDRACSNMPTrap.sh 1/agentx -force

- If -force is not specified, ensure that the net-SNMP is configured and restart the snmpd service.
- To enable this feature:

```
Enable-iDRACSNMPTrap.sh 1
```

```
Enable-iDRACSNMPTrap.sh enable
```

• To disable this feature:

```
Enable-iDRACSNMPTrap.sh 0
```

Enable-iDRACSNMPTrap.sh disable

NOTE: The **--force** option configures the Net-SNMP to forward the traps. However, you must configure the trap destination.

VMware ESXi operating system

On all iSM supported ESXi operating systems, the iSM v2.3 supports a Common Management Programming Interface (CMPI) method provider to enable this feature remotely by using the WinRM remote commands.

```
winrm i EnableInBandSNMPTraps http://schemas.dell.com/wbem/
wscim/1/cim-schema/2/root/cimv2/dcim/DCIM_iSMService? __cimnamespace=root/cimv2/
dcim+InstanceID=iSMExportedFunctions -u:<user-name> -p:<passwd> -r:https://<remote-host-
name</pre>
```

ip-address>:443/WSMan -a:basic -encoding:utf-8 -skipCNCheck -skipCACheck -skipRevocationcheck @{state="[0/1]"}

(i) NOTE: You must review and configure the VMware ESXi system-wide SNMP settings for traps.

NOTE: For more details, refer to the **In-BandSNMPAlerts** technical white paper available at https://www.dell.com/ support.

iDRAC access via Host OS

By using this feature, you can configure and monitor the hardware parameters through iDRAC Web interface, WSMan, and RedFish interfaces using the host IP address without configuring the iDRAC IP address. You can use the default iDRAC credentials if the iDRAC server is not configured or continue to use the same iDRAC credentials if the iDRAC server was configured earlier.

iDRAC access via Windows Operating Systems

You can perform this task by using the following methods:

- Install the iDRAC access feature by using the webpack.
- Configure using iSM PowerShell script

Installation by using MSI

292 Using iDRAC Service Module



You can install this feature by using the web-pack. This feature is disabled on a typical iSM installation. If enabled, the default listening port number is 1266. You can modify this port number within the range 1024 through 65535. iSM redirects the connection to the iDRAC. iSM then creates an inbound firewall rule, OS2iDRAC. The listening port number is added to the OS2iDRAC firewall rule in the host operating system, which allows incoming connections. The firewall rule is enabled automatically when this feature is enabled.

Beginning with iSM 2.4.0, you can retrieve the current status and listening-port configuration by using the following Powershell cmdlet:

Enable-iDRACAccessHostRoute -status get

The output of this command indicates whether this feature is enabled or disabled. If the feature is enabled, it displays the listening-port number.

(i) NOTE: Ensure that the Microsoft IP Helper Services is running on your system for this feature to function.

To access the iDRAC Web interface, use the format https://<host-name>or OS-IP>:443/login.html in the browser, where:

- <host-name> Complete host name of the server on which iSM is installed and configured for iDRAC access via OS feature. You can use the OS IP address if the host name is not present.
- 443 Default iDRAC port number. This is called the Connect Port number to which all the incoming connections on listen port number are redirected. You can modify the port number through iDRAC Web interface, WSMan, and RACADM interfaces.

Configuration by using iSM PowerShell cmdlet

If this feature is disabled while installing iSM, you can enable the feature by using the following Windows PowerShell command provided by iSM:

Enable-iDRACAccessHostRoute

If the feature is already configured, you can disable or modify it by using the PowerShell command and the corresponding options. The available options are as follows:

- Status This parameter is mandatory. The values are not case sensitive and the value can be true, false, or get.
- **Port** This is the listening port number. If you do not provide a port number, the default port number (1266) is used. If the **Status** parameter value is FALSE, then you can ignore rest of the parameters. You must enter a new port number that is not already configured for this feature. The new port number settings overwrite the existing OS2iDRAC in-bound firewall rule and you can use the new port number to connect to iDRAC. The value range is from 1024 to 65535.
- **IPRange** This parameter is optional and it provides a range of IP addresses that are allowed to connect to iDRAC through the host operating system. The IP address range format is in Classless Inter-Domain Routing (CIDR) format, which is a combination of IP address and subnet mask. For example, 10.94.111.21/24. Access to iDRAC is restricted for IP addresses that are not within the range.

(i) NOTE: This feature supports only IPv4 addresses.

iDRAC access via Linux Operating Systems

You can install this feature by using the setup.sh file that is available with the web-pack. This feature is disabled on a default or typical iSM installation. To get the status of this feature, use the following command:

Enable-iDRACAccessHostRoute get-status

To install, enable, and configure this feature, use the following command:

./Enable-iDRACAccessHostRoute <Enable-Flag> [<source-port> <source-IP-range/source-iprange-mask>]

<Enable-Flag>=0

Disable

<source-port> and <source-IP-range/source-ip-range-mask> are not required.

<Enable-Flag>=1

Enable

<source-port> is required and <source-ip-range-mask> is optional.

<source-IP-range>

Using iDRAC Service Module 293



Coexistence of OpenManage Server Administrator and iDRAC Service Module

In a system, both OpenManage Server Administrator and the iDRAC Service Module can co-exist and continue to function correctly and independently.

If you have enabled the monitoring features during the iDRAC Service Module installation, then after the installation is complete if the iDRAC Service Module detects the presence of OpenManage Server Administrator, it disables the set of monitoring features that overlap. If OpenManage Server Administrator is running, the iDRAC Service Module disables the overlapping monitoring features after logging to the OS and iDRAC.

When you re-enable these monitoring features through the iDRAC interfaces later, the same checks are performed and the features are enabled depending on whether OpenManage Server Administrator is running or not.

Using iDRAC Service Module from iDRAC web interface

To use the iDRAC Service Module from the iDRAC web interface:

- Go to IDRAC Settings > Overview > iDRAC Service Module > Configure Service Module. The iDRAC Service Module Setup page is displayed.
- 2. You can view the following:
 - Installed iDRAC Service Module version on the host operating system
 - Connection status of the iDRAC Service Module with iDRAC.
 - () NOTE: When a server has multiple operating systems and iDRAC Service Module is installed in all operating systems, then iDRAC connects only with the most recent instance of iSM among all operating systems. An error is displayed for all the older instances of iSM on other operating systems. To connect iSM with iDRAC on any other operating system which already has iSM installed, uninstall and reinstall iSM on that particular operating system.
- 3. To perform out-of-band monitoring functions, select one or more of the following options:
 - **OS Information**—View the operating system information.
 - **Replicate Lifecycle Log in OS Log**—Include Lifecycle Controller logs to operating system logs. This option is disabled if OpenManage Server Administrator is installed on the system.
 - **WMI Information** Include WMI information.
 - Auto System Recovery Action—Perform auto recovery operations on the system after a specified time (in seconds):
 - Reboot
 - Power Off System
 - Power Cycle System

This option is disabled if OpenManage Server Administrator is installed on the system.

Using iDRAC Service Module from RACADM

To use the iDRAC Service Module from RACADM, use the objects in the ServiceModule group.

For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.



Using USB port for server management

On the 14th generation servers, a dedicated micro USB port is available to configure iDRAC. You can perform the following functions using the micro USB port:

- Connect to the system using the USB network interface to access system management tools such as iDRAC web interface and RACADM.
- Configure a server by using SCP files that are stored on a USB drive.
- (i) NOTE: To manage a USB port or to configure a server by importing Server Configuration Profile (SCP) files on a USB drive, you must have the System Control privilege.

(i) NOTE: An alert / report is generated when a USB device is inserted. This feature is only available on Intel based servers.

To configure Management USB Settings, go to **iDRAC Settings** > **Settings** > **Management USB Settings**. Following options are available:

• **USB Management Port**—Select **Enabled** to enable the port either to import the SCP file when a USB drive is connected or to access iDRAC using the micro USB port.

(i) NOTE: Ensure that the USB drive contains a valid SCP file.

(i) NOTE: Use an OTG adapter to convert from Type-A to Micro-B USB. Connections from USB hubs are not supported.

- **iDRAC Managed: USB SCP**—Select from following options to configure the system by importing SCP stored on a USB drive:
 - Disabled—Disables SCP imports
 - **Enabled only when server has default credential settings** If this option is selected then the SCP can only be imported when the default password is not changed for the following:
 - BIOS
 - iDRAC web interface
 - Enabled only for compressed configuration files—Select this option to allow SCP file import only if the files are in compressed format.

NOTE: Selecting this option allows you to password protect the compressed file. You can enter a password to secure the file by using **Password for Zip file** option.

• Enabled—Select this option to allow importing SCP file without running a check during runtime.

Topics:

- Accessing iDRAC interface over direct USB connection
- Configuring iDRAC using server configuration profile on USB device

Accessing iDRAC interface over direct USB connection

The iDRAC direct feature allows you to directly connect your laptop to the iDRAC USB port. This feature allows you to interact directly with the iDRAC interfaces such as the web interface, RACADM, and WSMan for advanced server management and servicing.

For a list of supported browsers and operating systems, see the *iDRAC Release Notes* available at https://www.dell.com/ idracmanuals.

(i) NOTE: If you are using Windows operating system, you may need to install an RNDIS driver to use this feature.

To access the iDRAC interface over the USB port:

- 1. Turn off any wireless networks and disconnect from any other hard wired network.
- 2. Ensure that the USB port is enabled. For more information, see Configuring USB management port settings.

Using USB port for server management 295



- **3.** Wait for the laptop to acquire IP address 169.254.0.4. It may take several seconds for the IP addresses to be acquired. iDRAC acquires the IP address 169.254.0.3.
- 4. Start using iDRAC network interfaces such as the web interface, RACADM, Redfish or WSMan. For example, to access the iDRAC web interface, open a supported browser, and type the address 169.254.0.3 and press enter.
- 5. When iDRAC is using the USB port, the LED blinks indicating activity. The blink frequency is four per second.
- **6.** After completing the desired actions, disconnect the USB cable from the system. The LED turns off.

Configuring iDRAC using server configuration profile on USB device

With the iDRAC USB management port, you can configure iDRAC at-the-server. Configure the USB Management port settings in iDRAC, insert the USB device that has the server configuration profile, and then import the server configuration profile from the USB device to iDRAC.

NOTE: You can set the USB Management port settings using the iDRAC interfaces only if there is no USB device connected to the server.

Configuring USB management port settings

You can enable or disable the iDRAC Direct USB Port using the system BIOS. Navigate to **System BIOS** > **Integrated Devices**. Select **On** to enable and **Off** to disable the iDRAC Direct USB Port.

In iDRAC, you must have Server Control privilege to configure the USB management port. When a USB device is connected, the **System Inventory** page displays the USB device information under the Hardware Inventory section.

An event is logged in the Lifecycle Controller logs when:

- The device is in Automatic or iDRAC mode, and USB device is inserted or removed.
- USB Management Port Mode is modified.
- Device is automatically switched from iDRAC to OS.
- Device is ejected from iDRAC or OS

When a device exceeds its power requirements as allowed by USB specification, the device is detached and an over-current event is generated with the following properties:

- Category : System Health
- Type: USB device
- Severity: Warning
- Allowed notifications: Email, SNMP trap, remote syslog, and WS-Eventing
- Actions: None

An error message is displayed and logged to Lifecycle Controller log when:

- You try to configure the USB management port without the Server Control user privilege.
- A USB device is in use by iDRAC, and you attempt to modify the USB Management Port Mode.
- A USB device is in use by iDRAC, and you remove the device.

Configuring USB management port using web interface

To configure the USB port:

- 1. In the iDRAC Web interface, go to **iDRAC Settings** > **Settings** > **Management USB Settings**.
- 2. The USB Management Port is set to Enabled.
- **3.** From the **iDRAC Managed: USB SCP** Configuration drop-down menu, select options to configure a server by importing Server Configuration Profile files stored on a USB drive:
 - Disabled
 - Enabled only when server has default credential settings
 - Enabled only for compressed configuration files
 - Enabled



For information about the fields, see the *iDRAC Online Help*.

() **NOTE:** iDRAC9 allows you to password protect the compressed file after you select Enabled only for compressed configuration files to compress the file before importing. You can enter a password to secure the file by using Password for Zip file option.

4. Click **Apply** to apply the settings.

Configuring USB management port using RACADM

To configure the USB management port, use the following RACADM sub commands and objects:

• To view the USB port status:

racadm get iDRAC.USB.PortStatus

• To view the USB port configuration:

racadm get iDRAC.USB.ManagementPortMode

• To view USB device inventory:

racadm hwinventory

• To set up overcurrent alert configuration:

```
racadm eventfilters
```

For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Configuring USB management port using iDRAC settings utility

To configure the USB port:

- In the iDRAC Settings Utility, go to Media and USB Port Settings. The iDRAC Settings Media and USB Port Settings page is displayed.
- 2. From the **iDRAC Direct: USB Configuration XML** drop-down menu, select options to configure a server by importing server configuration profile stored on a USB drive:
 - Disabled
 - Enabled while server has default credential settings only
 - Enabled only for compressed configuration files
 - Enabled
 - For information about the fields, see the *iDRAC Settings Utility Online Help*.
- 3. Click **Back**, click **Finish** and then click **Yes** to apply the settings.

Importing Server Configuration Profile from USB device

Make sure to create a directory in root of the USB device called System_Configuration_XML that contains both the config and control files:

- Server Configuration Profile (SCP) is in the System_Configuration_XML sub-directory under the USB device root directory. This file includes all the attribute-value pairs of the server. This includes attributes of iDRAC, PERC, RAID, and BIOS. You can edit this file to configure any attribute on the server. The file name can be <servicetag>-config.xml, <servicetag>-config.json, <modelnumber>-config.xml, <modelnumber>-config.json, config.yml or config.json.
- Control file Includes parameters to control the import operation and does not have attributes of iDRAC or any other component in the system. The control file contain three parameters:
 - ShutdownType Graceful, Forced, No Reboot.
 - TimeToWait (in secs) 300 minimum and 3600 maximum.
 - EndHostPowerState on/off.

Example of control.xml file:



```
<InstructionTable>
  <InstructionRow>
     <InstructionType>Configuration XML import Host control Instruction
     </InstructionType>
     <Instruction>ShutdownType</Instruction>
     <Value>NoReboot</Value>
     <ValuePossibilities>Graceful,Forced,NoReboot</ValuePossibilities>
  </InstructionRow>
  <InstructionRow>
     <InstructionType>Configuration XML import Host control Instruction
     </InstructionType>
     <Instruction>TimeToWait</Instruction>
     <Value>300</Value>
     <ValuePossibilities>Minimum value is 300 -Maximum value is
      3600 seconds.</ValuePossibilities>
  </InstructionRow>
  <InstructionRow>
     <InstructionType>Configuration XML import Host control Instruction
     </InstructionType>
     <Instruction>EndHostPowerState</Instruction>
     <Value>On</Value>
     <ValuePossibilities>On,Off</ValuePossibilities>
  </InstructionRow>
</InstructionTable>
```

You must have Server Control privilege to perform this operation.

NOTE: While importing the SCP, changing the USB management settings in the SCP file results in a failed job or job completed with errors. You can comment out the attributes in the SCP to avoid the errors.

To import the server configuration profile from the USB device to iDRAC:

1. Configure the USB management port:

- Set USB Management Port Mode to Automatic or iDRAC.
- Set iDRAC Managed: USB XML Configuration to Enabled with default credentials or Enabled.
- 2. Insert the USB key (that has the configuration.xml and the control.xml file) to the iDRAC USB port.

(i) NOTE: File name and file type are case sensitive for XML files. Ensure that both are in lower case.

- **3.** The server configuration profile is discovered on the USB device in the System_Configuration_XML sub-directory under the USB device root directory. It is discovered in the following sequence:
 - <servicetag>-config.xml/<servicetag>-config.json
 - <modelnum>-config.xml/<modelnum>-config.json
 - config.xml/config.json
- 4. A server configuration profile import job starts.

If the profile is not discovered, then the operation stops.

If **iDRAC Managed: USB XML Configuration** was set to **Enabled with default credentials** and the BIOS setup password is not null or if one of the iDRAC user accounts have been modified, an error message is displayed and the operation stops.

- 5. LCD panel and LED, if present, display the status that an import job has started.
- 6. If there is a configuration that needs to be staged and the **Shut Down Type** is specified as **No Reboot** is specified in the control file, you must reboot the server for the settings to be configured. Else, server is rebooted and the configuration is applied. Only when the server was already powered down, then the staged configuration is applied even if the **No Reboot** option is specified.
- 7. After the import job is complete, the LCD/LED indicates that the job is complete. If a reboot is required, LCD displays the job status as "Paused waiting on reboot".
- 8. If the USB device is left inserted on the server, the result of the import operation is recorded in the results.xml file in the USB device.



LCD messages

If the LCD panel is available, it displays the following messages in a sequence:

- 1. Importing When the server configuration profile is being copied from the USB device.
- 2. Applying When the job is in-progress.
- **3.** Completed When the job has completed successfully.
- 4. Completed with errors When the job has completed with errors.
- 5. Failed When the job has failed.

For more details, see the results file on the USB device.

LED blinking behavior

The USB LED indicates the status of a server-configuration profile operation being performed using the USB port. The LED may not be available on all systems.

- Solid green The server configuration profile is being copied from the USB device.
- Blinking green The job is in progress.
- Blinking amber The job has failed or completed with errors.
- Solid green The job has completed successfully.

NOTE: On PowerEdge R840 and R940xa, if there is a LCD present, the USB LED does not blink when an import operation is in progress using the USB port. Check the status of the operation using the LCD.

Logs and results file

The following information is logged for the import operation:

- Automatic import from USB is logged in the Lifecycle Controller log file.
- If the USB device is left inserted, the job results are recorded in the Results file located in the USB key.
- A Result file named Results.xml is updated or created in the subdirectory with the following information:
- Service tag Data is recorded after the import operation has either returned a job ID or returned an error.
- Job ID Data is recorded after the import operation has returned a job ID.
- Start Date and Time of Job Data is recorded after the import operation has returned a job ID.
- Status Data is recorded when the import operation returns an error or when the job results are available.



Using Quick Sync 2

With Dell OpenManage Mobile running on an Android or iOS mobile device, you can easily access server directly or through OpenManage Essentials or OpenManage Enterprise (OME) console. It allows you to review server details and inventory, view LC and System Event logs, get automatic notifications on mobile device from an OME console, assign IP address and modify iDRAC password, configure key BIOS attributes, and take remediation actions as needed. You can also power cycle a server, access system console, or access iDRAC GUI.

OMM can be downloaded for free from the Apple App Store, or from Google Play Store.

You must install the OpenManage Mobile application on the mobile device (supports Android 5.0+ and iOS 9.0+ mobile devices) to manage server using iDRAC Quick Sync 2 interface.

(i) NOTE: This section is displayed only in those servers that has Quick Sync 2 module in left rack ear.

(i) **NOTE:** This feature is currently supported on mobile devices with Android operating system and Apple iOS.

In the current release, this feature is available on all 14th generation of PowerEdge servers. It requires Quick Sync 2 Left Control Panel (embedded in **Left rack ear**) and Bluetooth Low Energy (and optionally Wi-Fi) enabled mobile devices. Therefore, it is a hardware up-sell and the feature capabilities are not dependent on iDRAC software licensing.

NOTE: For more information on configuring Quick Sync 2 in MX platform systems, see the OpenManage Enterprise Modular User's Guide and OpenManage Mobile User's Guide available at **dell.com/support/manuals**.

The iDRAC Quick Sync 2 Configuration procedures:

(i) NOTE: Not applicable for MX platforms.

Once Quick Sync is configured, activate the Quick Sync 2 button on the Left Control Panel. Make sure the Quick Sync 2 light turns on. Access the Quick Sync 2 Information using a mobile device (Android 5.0+ or IOS 9.0+, OMM 2.0 or above).

Using OpenManage Mobile, you can:

- View inventory information
- View monitoring information
- Configure the basic iDRAC network settings

For more information about OpenManage Mobile, see the *Dell EMC OpenManage Mobile User's Guide* available at https://www.dell.com/openmanagemanuals.

Topics:

- Configuring iDRAC Quick Sync 2
- Using mobile device to view iDRAC information

Configuring iDRAC Quick Sync 2

Using iDRAC web interface, RACADM, WSMan and iDRAC HII you can configure iDRAC Quick Sync 2 feature to allow access to the mobile device:

- Access Configure to Read-Write, Read-only, and Disabled. Read-Write is the default option.
- Timeout Configure to Enabled or Disabled. Enabled is the default option.
- **Timeout Limit** Indicates the time after which the Quick Sync 2 mode is disabled. By default, seconds are selected. The default value is 120 seconds. The range is 120 to 3600 seconds.
 - 1. If enabled, you can specify a time after which the Quick Sync 2 mode is turned off. To turn on, press the activation button again.
 - 2. If disabled, the timer does not allow you to enter a time-out period.
- **Read Authentication** Configures to Enabled, this is the default option.
- WiFi Configures to Enabled, this is the default option.



You must have Server Control privilege to configure the settings. A server reboot is not required for the settings to take effect. once configured, you can activate the Quick Sync 2 button on the Left Control Panel. Make sure the Quick Sync light turns on. Then, access the Quick Sync Information via a mobile device.

An entry is logged to the Lifecycle Controller log when the configuration is modified.

Configuring iDRAC Quick Sync 2 settings using web interface

To configure iDRAC Quick Sync 2:

- 1. In the iDRAC web interface, go to Configuration > System Settings > Hardware Settings > iDRAC Quick Sync.
- 2. In the **iDRAC Quick Sync** section, from the **Access** menu, select one of the following to provide access to the Android or iOS mobile device:
 - Read-write
 - Read-only
 - Disabled
- 3. Enable the Timer.
- **4.** Specify the Timeout Limit.
 - For more information about the fields, see the *iDRAC Online Help*.
- 5. Click **Apply** to apply the settings.

Configuring iDRAC Quick Sync 2 settings using RACADM

To configure the iDRAC Quick Sync 2feature, use the racadm objects in the **System.QuickSync** group. For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Configuring iDRAC Quick Sync 2 settings using iDRAC settings utility

To configure iDRAC Quick Sync 2:

- 1. In the iDRAC GUI, go to Configuration > Systems Settings > Hardware Settings > iDRAC Quick Sync.
- 2. In the iDRAC Quick Sync section:
 - Specify the access level.
 - Enable Timeout.
 - Specify the User Defined Timeout Limit (the range is 120 to 3600 seconds.).

For more information about the fields, see the *iDRAC Online Help*.

3. Click **Back**, click **Finish**, and then click **Yes**. The settings are applied.

Using mobile device to view iDRAC information

To view iDRAC information from the mobile device, see the *Dell EMC OpenManage Mobile User's Guide* available at https://www.dell.com/openmanagemanuals for the steps.



Managing virtual media

iDRAC provides virtual media with HTML5 based client with local ISO and IMG file, remote ISO and IMG file support. Virtual media allows the managed server to access media devices on the management station or ISO CD/DVD images on a network share as if they were devices on the managed server. You need iDRAC Configure privilege to modify the configuration.

Following are the configurable attributes:

- Attached Media Enabled Enabled / Disabled
- Attach Mode Auto-attach, Attached, and Detached
- Max Sessions 1
- Active Sessions 1
- Virtual Media Encryption Enabled (by default)
- Floppy Emulation Disabled (by default)
- Boot Once Enabled / Disabled
- Connection Status Connected / Disconnected

Using the Virtual Media feature, you can:

- Remotely access media connected to a remote system over the network
- Install applications
- Update drivers
- Install an operating system on the managed system

This is a licensed feature for rack and tower servers. It is available by default for blade servers.

The key features are:

- Virtual Media supports virtual optical drives (CD/DVD) and USB flash drives.
- You can attach only one USB flash drive, image, or key and one optical drive on the management station to a managed system. Supported optical drives include a maximum of one available optical drive or one ISO image file.

The following figure shows a typical Virtual Media setup.

- Any connected Virtual Media emulates a physical device on the managed system.
- On Windows-based managed systems, the Virtual Media drives are auto-mounted if they are attached and configured with a drive letter.
- On Linux-based managed systems with some configurations, the Virtual Media drives are not auto-mounted. To manually mount the drives, use the mount command.
- All the virtual drive access requests from the managed system are directed to the management station across the network.
- Virtual devices appear as two drives on the managed system without the media being installed in the drives.
- You can share the management station CD/DVD drive (read only), but not a USB media, between two managed systems.
- Virtual media requires a minimum available network bandwidth of 128 Kbps.
- If LOM or NIC failover occurs, then the Virtual Media session may be disconnected.

After attaching a Virtual Media image through Virtual Console, the drive may not show up in Windows host OS. Check Windows Device Manager for any unknown mass storage devices. Right click on the unknown device and update the driver or choose uninstall driver. The device is recognized by Windows after disconnecting and reconnecting vMedia.



Figure 4. Virtual media setup

Topics:

302 Managing virtual media



- Supported drives and devices
- Configuring virtual media
- Accessing virtual media
- Setting boot order through BIOS
- Enabling boot once for virtual media

Supported drives and devices

The following table lists the drives supported through virtual media.

Table 60. Supported drives and devices

Drive	Supported Storage Media
Virtual Optical Drives	 CD-ROM DVD CD-RW Combination drive with CD-ROM media
USB flash drives	 USB CD-ROM drive with CD-ROM media USB Key image in the ISO9660 format

Configuring virtual media

Before you configure the Virtual Media settings, make sure that you have configured your Web browser to use Java or ActiveX plug-in.

Configuring virtual media using iDRAC web interface

To configure virtual media settings:

- CAUTION: Do not reset iDRAC when running a Virtual Media session. Otherwise, undesirable results may occur, including data loss.
- 1. In the iDRAC Web interface, go to Configuration > Virtual Media > Attached Media.
- 2. Specify the required settings. For more information, see the *iDRAC Online Help*.
- 3. Click Apply to save the settings.

Configuring virtual media using RACADM

To configure the virtual media, use the set command with the objects in the **iDRAC.VirtualMedia** group.

For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Configuring virtual media using iDRAC settings utility

You can attach, detach, or auto-attach virtual media using the iDRAC Settings utility. To do this:

- 1. In the iDRAC Settings utility, go to Media and USB Port Settings. The iDRAC Settings Media and USB Port Settings page is displayed.
- 2. In the Virtual Media section, select Detach, Attach, or Auto attach based on the requirement. For more information about the options, see *iDRAC Settings Utility Online Help*.
- **3.** Click **Back**, click **Finish**, and then click **Yes**. The Virtual Media settings are configured.



Attached media state and system response

The following table describes the system response based on the Attached Media setting.

Table 61. Attached media state and system response

Attached Media State	System Response
Detach	Cannot map an image to the system.
Attach	Media is mapped even when Client View is closed.
Auto-attach	Media is mapped when Client View is opened and unmapped when Client View is closed.

Server settings for viewing virtual devices in virtual media

You must configure the following settings in the management station to allow visibility of empty drives. To do this, in Windows Explorer, from the **Organize** menu, click **Folder and search options**. On the **View** tab, deselect **Hide empty drives in the Computer folder** option and click **OK**.

Accessing virtual media

You can access Virtual Media with or without using the Virtual Console. Before you access Virtual Media, make sure to configure your Web browser(s).

Virtual Media and RFS are mutually exclusive. If the RFS connection is active and you attempt to launch the Virtual Media client, the following error message is displayed: *Virtual Media is currently unavailable. A Virtual Media or Remote File Share session is in use.*

If the RFS connection is not active and you attempt to launch the Virtual Media client, the client launches successfully. You can then use the Virtual Media client to map devices and files to the Virtual Media virtual drives.

Launching virtual media using virtual console

Before you launch Virtual Media through the Virtual Console, make sure that:

- Virtual Console is enabled.
- System is configured to not hide empty drives In Windows Explorer, navigate to Folder Options, clear the Hide empty drives in the Computer folder option, and click OK.

To access Virtual Media using Virtual Console:

- In the iDRAC web interface, go to Configuration > Virtual Console. The Virtual Console page is displayed.
- 2. Click Launch Virtual Console. The Virtual Console Viewer is launched.

i NOTE: On Linux, Java is the default plug-in type for accessing the Virtual Console. On Windows, open the .jnlp file to

- launch the Virtual Console using Java.
- 3. Click Virtual Media > Connect Virtual Media.

The Virtual Media session is established and the Virtual Media menu displays the list of devices available for mapping.

(i) NOTE: The Virtual Console Viewer window must remain active while you access the Virtual Media.

Launching virtual media without using virtual console

Before you launch Virtual Media when the **Virtual Console** is disabled, ensure that System is configured to unhide empty drives. To do this, in Windows Explorer, go to **Folder Options**, clear the **Hide empty drives in the Computer folder** option, and click **OK**.

To access Virtual Media when Virtual Console is disabled:

304 Managing virtual media



- 1. In the iDRAC web Interface, go to **Configuration** > Virtual Media.
- 2. Click Connect Virtual Media.

Alternatively, you can also launch the Virtual Media by following these steps:

- 1. Go to Configuration > Virtual Console.
- 2. Click Launch Virtual Console. The following message is displayed:

Virtual Console has been disabled. Do you want to continue using Virtual Media redirection?

- 3. Click OK. The Virtual Media window is displayed.
- 4. From the Virtual Media menu, click Map CD/DVD or Map Removable Disk. For more information, see Mapping virtual drive.
- 5. Virtual Media Statistics shows the list of target drives, their mapping, status (Read-Only or Not), Duration of connection, Read/Write Bytes, and the transfer rate.
- (i) **NOTE:** The virtual device drive letters on the managed system do not coincide with the physical drive letters on the management station.

() NOTE: The Virtual Media may not function correctly on systems running Windows operating system configured with Internet Explorer Enhanced Security. To resolve this issue, see the Microsoft operating system documentation or contact the system administrator.

Adding virtual media images

You can create a media image of the remote folder and mount it as a USB attached device to the server's operating system. To add Virtual Media images:

- 1. Click Virtual Media > Create Image....
- 2. In the **Source Folder** field, click **Browse** and browse to the folder or directory to be used as the source for the image file. The image file is on the management station or the C: drive of the managed system.
- 3. In the **Image File Name** field, the default path to store the created image files (typically the desktop directory) appears. To change this location, click **Browse** and navigate to a location.
- 4. Click Create Image.

The image creation process starts. If the image file location is within the source folder, a warning message is displayed indicating that the image creation cannot proceed as the image file location within the source folder causes an infinite loop. If the image file location is not within the source folder, then the image creation proceeds.

After the image is created, a success message is displayed.

5. Click Finish.

The image is created.

When a folder is added as an image, a **.img** file is created on the Desktop of the management station from which this feature is used. If this **.img** file is moved or deleted, then the corresponding entry for this folder in the **Virtual Media** menu does not work. Therefore, it is recommended not to move or delete the **.img** file while the *image* is being used. However, the **.img** file can be removed after the relevant entry is first deselected and then removed using **Remove Image** to remove the entry.

Viewing virtual device details

To view the virtual device details, in the Virtual Console Viewer, click **Tools** > **Stats**. In the **Stats** window, the **Virtual Media** section displays the mapped virtual devices and the read/write activity for each device. If Virtual Media is connected, this information is displayed. If Virtual Media is not connected, the "Virtual Media is not connected" message is displayed.

If the Virtual Media is launched without using the Virtual Console, then the **Virtual Media** section is displayed as a dialog box. It provides information about the mapped devices.

Accessing drivers

Dell EMC PowerEdge servers have all the supported operating system drivers embedded on the system flash memory. Using iDRAC, you can mount or unmount drivers easily to deploy the operating system on your server.

To mount the drivers:



- 1. On the iDRAC web interface, go to Configuration > Virtual Media.
- 2. Click Mount Drivers.
- 3. Select the OS from the pop-up window and click Mount Drivers.

(i) NOTE: The Expose duration is 18 hours by default.

To unmount the drivers post completion of the mount:

- 1. Go to Configuration > Virtual Media.
- 2. Click Unmount Drivers.
- **3.** Click \mathbf{OK} on the pop-up window.
- (i) **NOTE:** The **Mount Drivers** option may not be displayed if the driver pack is not available on the system. Ensure to download and install the latest driver pack from https://www.dell.com/support.

Resetting USB

To reset the USB device:

- In the Virtual Console viewer, click Tools > Stats. The Stats window is displayed.
- Under Virtual Media, click USB Reset.
 A message is displayed warning the user that resetting the USB connection can affect all the input to the target device including Virtual Media, keyboard, and mouse.
- 3. Click Yes.
 - The USB is reset.

(i) NOTE: iDRAC Virtual Media does not terminate even after you log out of iDRAC Web interface session.

Mapping virtual drive

To map the virtual drive:

- **NOTE:** While using ActiveX or Java-based Virtual Media, you must have administrative privileges to map an operating system DVD or a USB flash drive (that is connected to the management station). To map the drives, launch IE as an administrator or add the iDRAC IP address to the list of trusted sites.
- 1. To establish a Virtual Media session, from the Virtual Media menu, click Connect Virtual Media.

For each device available for mapping from the host server, a menu item appears under the **Virtual Media** menu. The menu item is named according to the device type such as:

- Map CD/DVD
- Map Removable Disk

The Map DVD/CD option can be used for ISO files and the Map Removable Disk option can be used for images.

() NOTE:

- You cannot map physical media such USB-based drives, CD, or DVD by using the HTML5 based virtual console.
- You cannot map USB keys as virtual media disks using Virtual Console/Virtual media over a RDP session.
- You cannot map physical media with NTFS format in ehtml removable media, use FAT or exFAT devices

2. Click the device type that you want to map.

NOTE: The active session displays if a Virtual Media session is currently active from the current Web interface session, from another Web interface session.

3. In the Drive/Image File field, select the device from the drop-down list.

The list contains all the available (unmapped) devices that you can map (CD/DVD and Removable Disk) and image file types that you can map (ISO or IMG). The image files are located in the default image file directory (typically the user's desktop). If the device is not available in the drop-down list, click **Browse** to specify the device.

The correct file type for CD/DVD is ISO and for removable disk it is IMG.

If the image is created in the default path (Desktop), when you select **Map Removable Disk**, the created image is available for selection in the drop-down menu.



If image is created in a different location, when you select **Map Removable Disk**, the created image is not available for selection in the drop-down menu. Click **Browse** to specify the image.

() NOTE:

- **Read only** option will be grayed out in ehtml5 based JAVA removable media.
- Floppy emulation is not supported in ehtml5 plugin.
- 4. Select Read-only to map writable devices as read-only.

For CD/DVD devices, this option is enabled by default and you cannot disable it.

(i) NOTE: The ISO and IMG files map as read-only files if you map these files by using the HTML5 virtual console.

5. Click Map Device to map the device to the host server.

After the device/file is mapped, the name of its **Virtual Media** menu item changes to indicate the device name. For example, if the CD/DVD device is mapped to an image file named foo.iso, then the CD/DVD menu item on the Virtual Media menu is named **foo.iso mapped to CD/DVD**. A check mark for that menu item indicates that it is mapped.

Displaying correct virtual drives for mapping

On a Linux-based management station, the Virtual Media **Client** window may display removable disks that are not part of the management station. To make sure that the correct virtual drives are available to map, you must enable the port setting for the connected SATA hard drive. To do this:

- 1. Reboot the operating system on the management station. During POST, press <F2> to enter System Setup.
- 2. Go to SATA settings. The port details are displayed.
- 3. Enable the ports that are actually present and connected to the hard drive.
- 4. Access the Virtual Media **Client** window. It displays the correct drives that can be mapped.

Unmapping virtual drive

To unmap the virtual drive:

- 1. From the Virtual Media menu, do any of the following:
 - Click the device that you want to unmap.
 - Click Disconnect Virtual Media.

A message appears asking for confirmation.

- 2. Click Yes.
 - The check mark for that menu item does not appear indicating that it is not mapped to the host server.
 - () NOTE: After unmapping a USB device attached to vKVM from a client system running the Macintosh operating system, the unmapped device may be unavailable on the client. Restart the system or manually mount the device on the client system to view the device.

(i) NOTE: To unmap a virtual DVD drive on Linux OS, unmount the drive and eject it.

Setting boot order through BIOS

Using the System BIOS Settings utility, you can set the managed system to boot from virtual optical drives or virtual floppy drives.

(i) NOTE: Changing Virtual Media while connected may stop the system boot sequence.

To enable the managed system to boot:

- 1. Boot the managed system.
- 2. Press <F2> to enter the System Setup page.
- 3. Go to System BIOS Settings > Boot Settings > BIOS Boot Settings > Boot Sequence.

In the pop-up window, the virtual optical drives and virtual floppy drives are listed with the standard boot devices.

4. Make sure that the virtual drive is enabled and listed as the first device with bootable media. If required, follow the on-screen instructions to modify the boot order.

Managing virtual media 307



- 5. Click OK, navigate back to System BIOS Settings page, and click Finish.
- 6. Click Yes to save the changes and exit.

The managed system reboots.

The managed system attempts to boot from a bootable device based on the boot order. If the virtual device is connected and a bootable media is present, the system boots to the virtual device. Otherwise, the system overlooks the device—similar to a physical device without bootable media.

Enabling boot once for virtual media

You can change the boot order only once when you boot after attaching remote Virtual Media device.

- Before you enable the boot once option, make sure that:
- You have *Configure User* privilege.
- Map the local or virtual drives (CD/DVD, Floppy, or USB flash device) with the bootable media or image using the Virtual Media options
- Virtual Media is in Attached state for the virtual drives to appear in the boot sequence.
- To enable the boot once option and boot the managed system from the Virtual Media:
- 1. In the iDRAC Web interface, go to **Overview** > **Server** > **Attached Media**.
- 2. Under Virtual Media, select the Enable Boot Once and click Apply.
- 3. Turn on the managed system and press <F2> during boot.
- 4. Change the boot sequence to boot from the remote Virtual Media device.
- 5. Reboot the server.
- The managed system boots once from the Virtual Media.


Managing vFlash SD card

(i) NOTE: vFlash is supported on AMD platform servers.

The vFlash SD card is a Secure Digital (SD) card that can be ordered and installed from the factory. You can use a card with a maximum of 16 GB capacity. After you insert the card, you must enable vFlash functionality to create and manage partitions. vFlash is a licensed feature.

(i) **NOTE:** There is no limitation of the size of SD card, you can open and replace the factory installed SD card with a higher capacity SD card. Since vFlash uses FAT32 file system, file size is limited to 4GB.

If the card is not available in the system's vFlash SD card slot, the following error message is displayed in the iDRAC Web interface at **Overview** > **Server** > **vFlash**:

SD card not detected. Please insert an SD card of size 256MB or greater.

() **NOTE:** Make sure that you only insert a vFlash compatible SD card in the iDRAC vFlash card slot. If you insert a noncompatible SD card, the following error message is displayed when you initialize the card: *An error has occurred while initializing SD card.*

The key features are:

- Provides storage space and emulates USB device (s).
- Create up to 16 partitions. These partitions, when attached, are exposed to the system as a Floppy drive, Hard Disk drive, or a CD/DVD drive depending on the selected emulation mode.
- Create partitions from supported file system types. Supports .img format for floppy, .iso format for CD/DVD, and both .iso and .img formats for Hard Disk emulation types.
- Create bootable USB device(s).
- Boot once to an emulated USB device.

NOTE: It is possible that a vFlash license may expire during a vFlash operation. If it happens, the on-going vFlash operations complete normally.

(i) NOTE: If FIPS mode is enabled, you cannot perform any vFlash actions.

Topics:

- Configuring vFlash SD card
- Managing vFlash partitions

Configuring vFlash SD card

Before configuring vFlash, make sure that the vFlash SD card is installed on the system. For information on how to install and remove the card from your system, see the *Installation and Service Manual* available at https://www.dell.com/poweredgemanuals.

(i) NOTE: You must have Access Virtual Media privilege to enable or disable vFlash functionality, and initialize the card.

Viewing vFlash SD card properties

After vFlash functionality is enabled, you can view the SD card properties using iDRAC Web interface or RACADM.



Viewing vFlash SD card properties using web interface

To view the vFlash SD card properties, in the iDRAC Web interface, go to **Configuration** > **System Settings** > **Hardware Settings** > **vFlash**. The Card Properties page is displayed. For information about the displayed properties, see the *iDRAC Online Help*.

Viewing vFlash SD card properties using RACADM

To view the vFlash SD card properties using RACADM, use the get command with the following objects:

- iDRAC.vflashsd.AvailableSize
- iDRAC.vflashsd.Health
- iDRAC.vflashsd.Licensed
- iDRAC.vflashsd.Size
- iDRAC.vflashsd.WriteProtect

For more information about these objects, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Viewing vFlash SD card properties using iDRAC settings utility

To view the vFlash SD card properties, in the **iDRAC Settings Utility**, go to **Media and USB Port Settings**. The **Media and USB Port Settings** page displays the properties. For information about the displayed properties, see the *iDRAC Settings Utility Online Help*.

Enabling or disabling vFlash functionality

You must enable the vFlash functionality to perform partition management.

Enabling or disabling vFlash functionality using web interface

To enable or disable the vFlash functionality:

- 1. In the iDRAC web interface, go to **Configuration** > **System Settings** > **Hardware Settings** > **vFlash**. The **SD Card Properties** page is displayed.
- 2. Select or clear the **vFLASH Enabled** option to enable or disable the vFlash functionality. If any vFlash partition is attached, you cannot disable vFlash and an error message is displayed.

(i) NOTE: If vFlash functionality is disabled, SD card properties are not displayed.

3. Click Apply. The vFlash functionality is enabled or disabled based on the selection.

Enabling or disabling vFlash functionality using RACADM

To enable or disable the vFlash functionality using RACADM:

racadm set iDRAC.vflashsd.Enable [n]

n=0

Disabled

n=1

Enabled

NOTE: The RACADM command functions only if a vFlash SD card is present. If a card is not present, the following message is displayed: *ERROR: SD Card not present*.



Enabling or disabling vFlash functionality using iDRAC settings utility

To enable or disable the vFlash functionality:

- 1. In the iDRAC Settings utility, go to Media and USB Port Settings. The iDRAC Settings . Media and USB Port Settings page is displayed.
- 2. In the **vFlash Media** section, select **Enabled** to enable vFlash functionality or select **Disabled** to disable the vFlash functionality.
- Click Back, click Finish, and then click Yes. The vFlash functionality is enabled or disabled based on the selection.

Initializing vFlash SD card

The initialize operation reformats the SD card and configures the initial vFlash system information on the card.

(i) **NOTE:** If the SD card is write-protected, then the Initialize option is disabled.

Initializing vFlash SD card using web interface

To initialize the vFlash SD card:

- In the iDRAC Web interface, go to Configuration > System Settings > Hardware Settings > vFlash. The SD Card Properties page is displayed.
- 2. Enable vFLASH and click Initialize.

All existing contents are removed and the card is reformatted with the new vFlash system information.

If any vFlash partition is attached, the initialize operation fails and an error message is displayed.

Initializing vFlash SD card using RACADM

To initialize the vFlash SD card using RACADM:

racadm set iDRAC.vflashsd.Initialized 1

All existing partitions are deleted and the card is reformatted.

For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Initializing vFlash SD card using iDRAC settings utility

To initialize the vFlash SD card using iDRAC Settings utility:

- 1. In the iDRAC Settings utility, go to **Media and USB Port Settings**.
- The iDRAC Settings . Media and USB Port Settings page is displayed.
- 2. Click Initialize vFlash.
- 3. Click Yes. The initialization operation starts.
- 4. Click **Back** and navigate to the same **iDRAC Settings**. **Media and USB Port Settings** page to view the successful message.

All existing contents are removed and the card is reformatted with the new vFlash system information.

Getting the last status using RACADM

To get the status of the last initialize command sent to the vFlash SD card:

- 1. Open SSH or Serial console to the system and log in.
- 2. Enter the command: racadm vFlashsd status The status of commands sent to the SD card is displayed.
- 3. To get the last status of all the vflash partitions, use the command:racadm vflashpartition status -a

Managing vFlash SD card 311



4. To get the last status of a particular partition, use command:racadm vflashpartition status -i (index)

(i) NOTE: If iDRAC is reset, the status of the last partition operation is lost.

Managing vFlash partitions

You can perform the following using the iDRAC Web interface or RACADM:

NOTE: An administrator can perform all operations on the vFlash partitions. Else, you must have **Access Virtual Media** privilege to create, delete, format, attach, detach, or copy the contents for the partition.

- Creating an empty partition
- Creating a partition using an image file
- Formatting a partition
- Viewing available partitions
- Modifying a partition
- Attaching or detaching partitions
- Deleting existing partitions
- Downloading partition contents
- Booting to a partition

() NOTE: If you click any option on the vFlash pages when an application such as WSMan, iDRAC Settings utility, or RACADM is using vFlash, or if you navigate to some other page in the GUI, iDRAC may display the message: vFlash is currently in use by another process. Try again after some time.

vFlash is capable of performing fast partition creation when there is no other on-going vFlash operation such as formatting, attaching partitions, and so on. Therefore, it is recommended to first create all partitions before performing other individual partition operations.

Creating an empty partition

An empty partition, when attached to the system, is similar to an empty USB flash drive. You can create empty partitions on a vFlash SD card. You can create partitions of type *Floppy* or *Hard Disk*. The partition type CD is supported only while creating partitions using images.

Before creating an empty partition, make sure that:

- You have Access Virtual Media privilege.
- The card is initialized.
- The card is not write-protected.
- An initialize operation is not being performed on the card.

Creating an empty partition using the web interface

To create an empty vFlash partition:

- 1. In iDRAC Web interface, go to Configuration > Systems Settings > Hardware Settings > vFlash > Create Empty Partition.
 - The Create Empty Partition page is displayed.
- Specify the required information and click Apply. For information about the options, see the *iDRAC Online Help*.
 A new unformatted empty partition is created that is read-only by default. A page indicating the progress percentage is displayed. An error message is displayed if:
 - The card is write-protected.
 - The label name matches the label of an existing partition.
 - A non-integer value is entered for the partition size, the value exceeds the available space on the card, or the partition size is greater than 4 GB.
 - An initialize operation is being performed on the card.



Creating an empty partition using RACADM

To create an empty partition:

- 1. Log in to the system using SSH or Serial console.
- 2. Enter the command:

racadm vflashpartition create -i 1 -o drive1 -t empty -e HDD -f fat16 -s [n]

where [n] is the partition size.

By default, an empty partition is created as read-write.

If the share is not configured using Username / Password, you need to specify the parameters as

```
-u anonymous -p anonymous
```

Creating a partition using an image file

You can create a new partition on the vFlash SD card using an image file (available in the **.img** or **.iso** format.) The partitions are of emulation types: Floppy (**.img**), Hard Disk (**.img**), or CD (**.iso**). The created partition size is equal to the image file size.

Before creating a partition from an image file, make sure that:

- You have Access Virtual Media privilege.
- The card is initialized.
- The card is not write-protected.
- An initialize operation is not being performed on the card.
- The image type and the emulation type match.
 - **NOTE:** The uploaded image and the emulation type must match. There are issues when iDRAC emulates a device with incorrect image type. For example, if the partition is created using an ISO image and the emulation type is specified as Hard Disk, then the BIOS cannot boot from this image.
- Image file size is less than or equal to the available space on the card.
- Image file size is less than or equal to 4 GB as the maximum partition size supported is 4 GB. However, while creating a partition using a Web browser, the image file size must be less than 2 GB.

(i) NOTE: The vFlash partition is an image file on a FAT32 file system. Thus, the image file has the 4 GB limitation.

(i) NOTE: Installation of a full OS is not supported.

Creating a partition using an image file using web interface

To create a vFlash partition from an image file:

- In iDRAC Web interface, go to Configuration > System Settings > Hardware Settings > vFlash > Create From Image. The Create Partition from Image File page is displayed.
- 2. Enter the required information and click **Apply**. For information about the options, see the *iDRAC Online Help*.

A new partition is created. For CD emulation type, a read-only partition is created. For Floppy or Hard Disk emulation type, a read-write partition is created. An error message is displayed if:

- The card is write-protected
- The label name matches the label of an existing partition.
- The size of the image file is greater than 4 GB or exceeds the available space on the card.
- The image file does not exist or the image file extension is neither .img nor .iso.
- An initialize operation is already being performed on the card.

Creating a partition from an image file using RACADM

To create a partition from an image file using RACADM:



- 1. Log in to the system using SSH or Serial console.
- 2. Enter the command

```
racadm vflashpartition create -i 1 -o drivel -e HDD -t image -l //myserver/sharedfolder/foo.iso -u root -p mypassword
```

By default, the created partition is read-only. This command is case sensitive for the image file name extension. If the file name extension is in upper case, for example FOO.ISO instead of FOO.iso, then the command returns a syntax error.

NOTE: Creating vFlash partition from an image file located on the CFS or NFS IPv6 enabled network share is not supported.

If the share is not configured using Username / Password, you need to specify the parameters as

```
-u anonymous -p anonymous
```

Formatting a partition

You can format an existing partition on the vFlash SD card based on the type of file system. The supported file system types are EXT2, EXT3, FAT16, and FAT32. You can only format partitions of type Hard Disk or Floppy, and not CD. You cannot format read-only partitions.

Before creating a partition from an image file, ensure that:

- You have Access Virtual Media privilege.
- The card is initialized.
- The card is not write-protected.
- An initialize operation is not being performed on the card.

To format vFlash partition:

- In iDRAC Web interface, go to Configuration > System Settings > Hardware Settings > vFlash > Format. The Format Partition page is displayed.
- 2. Enter the required information and click Apply.

For information about the options, see the *iDRAC Online Help*.

A warning message indicating that all the data on the partition will be erased is displayed.

3. Click OK.

The selected partition is formatted to the specified file system type. An error message is displayed if:

- The card is write-protected.
- An initialize operation is already being performed on the card.

Viewing available partitions

Make sure that the vFlash functionality is enabled to view the list of available partitions.

Viewing available partitions using web interface

To view the available vFlash partitions, in the iDRAC Web interface, go to **Configuration** > **System Settings** > **Hardware Settings** > **vFlash** > **Manage**. The **Manage Partitions** page is displayed listing the available partitions and related information for each partition. For information on the partitions, see the *iDRAC Online Help*.

Viewing available partitions using RACADM

To view the available partitions and their properties using RACADM:

1. Open a SSH or Serial console to the system and log in.



- 2. Enter the following commands:
 - To list all existing partitions and its properties: racadm vflashpartition list
 - To get the status of operation on partition 1: racadm vflashpartition status -i 1
 - To get the status of all existing partitions: racadm vflashpartition status -a

(i) NOTE: The -a option is valid only with the status action.

Modifying a partition

You can change a read-only partition to read-write or vice-versa. Before modifying the partition, make sure that:

- The vFlash functionality is enabled.
- You have Access Virtual Media privileges.

NOTE: By default, a read-only partition is created.

Modifying a partition using web interface

To modify a partition:

- 1. In the iDRAC Web interface, go to Configuration > System Settings > Hardware Settings > vFlash > Manage. The Manage Partitions page is displayed.
- 2. In the Read-Only column:
 - Select the checkbox for the partition(s) and click **Apply** to change to read-only.
 - Clear the checkbox for the partition(s) and click **Apply** to change to read-write.

The partitions are changed to read-only or read-write, based on the selections.

NOTE: If the partition is of type CD, the state is read-only. You cannot change the state to read-write. If the partition is attached, the check box is grayed-out.

Modifying a partition using RACADM

To view the available partitions and their properties on the card:

- 1. Log in to the system using SSH or Serial console.
- **2.** Use one of the following:
 - Using set command to change the read-write state of the partition:
 - To change a read-only partition to read-write:

racadm set iDRAC.vflashpartition.<index>.AccessType 1

• To change a read-write partition to read-only:

racadm set iDRAC.vflashpartition.<index>.AccessType 0

• Using set command to specify the Emulation type:

racadm set iDRAC.vflashpartition.<index>.EmulationType <HDD, Floppy, or CD-DVD>



Attaching or detaching partitions

When you attach one or more partitions, they are visible to the operating system and BIOS as USB mass storage devices. When you attach multiple partitions, based on the assigned index, they are listed in an ascending order in the operating system and the BIOS boot order menu.

If you detach a partition, it is not visible in the operating system and the BIOS boot order menu.

When you attach or detach a partition, the USB bus in the managed system is reset. This affects applications that are using vFlash and disconnects the iDRAC Virtual Media sessions.

Before attaching or detaching a partition, make sure that:

- The vFlash functionality is enabled.
- An initialize operation is not already being performed on the card.
- You have Access Virtual Media privileges.

Attaching or detaching partitions using web interface

To attach or detach partitions:

- In the iDRAC Web interface, go to Configuration > System Settings > Hardware Settings > vFlash > Manage. The Manage Partitions page is displayed.
- 2. In the Attached column:
 - Select the checkbox for the partition(s) and click **Apply** to attach the partition(s).
 - Clear the checkbox for the partition(s) and click **Apply** to detach the partition(s).

The partitions are attached or detached, based on the selections.

Attaching or detaching partitions using RACADM

To attach or detach partitions:

- 1. Log in to the system using SSH or Serial console.
- 2. Use the following commands:
 - To attach a partition:

```
racadm set iDRAC.vflashpartition.<index>.AttachState 1
```

• To detach a partition:

```
racadm set iDRAC.vflashpartition.<index>.AttachState 0
```

Operating system behavior for attached partitions

For Windows and Linux operating systems:

- The operating system controls and assigns the drive letters to the attached partitions.
- Read-only partitions are read-only drives in the operating system.
- The operating system must support the file system of an attached partition. Else, you cannot read or modify the contents of the partition from the operating system. For example, in a Windows environment the operating system cannot read the partition type EXT2 which is native to Linux. Also, in a Linux environment the operating system cannot read the partition type NTFS which is native to Windows.
- The vFlash partition label is different from the volume name of the file system on the emulated USB device. You can change the volume name of the emulated USB device from the operating system. However, it does not change the partition label name stored in iDRAC.

Deleting existing partitions

Before deleting existing partition(s), make sure that:



- The vFlash functionality is enabled.
- The card is not write-protected.
- The partition is not attached.
- An initialize operation is not being performed on the card.

Deleting existing partitions using web interface

To delete an existing partition:

- 1. In the iDRAC Web interface, go to Configuration > System Settings > Hardware Settings > vFlash > Manage. The Manage Partitions page is displayed.
- 2. In the **Delete** column, click the delete icon for the partition that you want to delete. A message is displayed indicating that this action permanently deletes the partition.
- **3.** Click **OK**. The partition is deleted.

Deleting existing partitions using RACADM

To delete partitions:

- 1. Open a SSH or Serial console to the system and log in.
- 2. Enter the following commands:
 - To delete a partition:

racadm vflashpartition delete -i 1

• To delete all partitions, re-initialize the vFlash SD card.

Downloading partition contents

You can download the contents of a vFlash partition in the .img or .iso format to the:

• Managed system (where iDRAC is operated from)

• Network location mapped to a management station.

Before downloading the partition contents, make sure that:

- You have Access Virtual Media privileges.
- The vFlash functionality is enabled.
- An initialize operation is not being performed on the card.
- For a read-write partition, it must not be attached.

To download the contents of the vFlash partition:

- In the iDRAC Web interface, go to Configuration > System Settings > Hardware Settings > vFlash > Download. The Download Partition page is displayed.
- 2. From the Label drop-down menu, select a partition that you want to download and click Download.
 - (i) **NOTE:** All existing partitions (except attached partitions) are displayed in the list. The first partition is selected by default.
- 3. Specify the location to save the file.

The contents of the selected partition are downloaded to the specified location.

NOTE: If only the folder location is specified, then the partition label is used as the file name, along with the

extension **.iso** for CD and Hard Disk type partitions, and **.img** for Floppy and Hard Disk type partitions.

Booting to a partition

You can set an attached vFlash partition as the boot device for the next boot operation.

Before booting a partition, make sure that:



- The vFlash partition contains a bootable image (in the .img or .iso format) to boot from the device.
- The vFlash functionality is enabled.
- You have Access Virtual Media privileges.

Booting to a partition using web interface

To set the vFlash partition as a first boot device, see Booting to a partition using web interface.

NOTE: If the attached vFlash partition(s) are not listed in the **First Boot Device** drop-down menu, make sure that the BIOS is updated to the latest version.

Booting to a partition using RACADM

To set a vFlash partition as the first boot device, use the iDRAC.ServerBoot object.

For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

() NOTE: When you run this command, the vFlash partition label is automatically set to boot once

(iDRAC.ServerBoot.BootOnce is set to 1.) Boot once boots the device to the partition only once and does not keep it persistently first in the boot order.



Using SMCLP

(i) NOTE: SMCLP is only supported in iDRAC versions earlier than 4.00.00.00.

The Server Management Command Line Protocol (SMCLP) specification enables CLI-based systems management. It defines a protocol for management commands transmitted over standard character oriented streams. This protocol accesses a Common Information Model Object Manager (CIMOM) using a human-oriented command set. The SMCLP is a sub-component of the Distributed Management Task Force (DMTF) SMASH initiative to streamline systems management across multiple platforms. The SMCLP specification, along with the Managed Element Addressing Specification and numerous profiles to SMCLP mapping specifications, describes the standard verbs and targets for various management task executions.

NOTE: It is assumed that you are familiar with the Systems Management Architecture for Server Hardware (SMASH) Initiative and the Server Management Working Group (SMWG) SMCLP specifications.

The SM-CLP is a subcomponent of the Distributed Management Task Force (DMTF) SMASH initiative to streamline server management across multiple platforms. The SM-CLP specification, along with the Managed Element Addressing Specification and numerous profiles to SM-CLP mapping specifications, describes the standard verbs and targets for various management task executions.

The SMCLP is hosted from the iDRAC controller firmware and supports SSH and serial-based interfaces. The iDRAC SMCLP interface is based on the SMCLP Specification Version 1.0 provided by the DMTF organization.

NOTE: Information about the profiles, extensions, and MOFs are available at https://www.dell.com/support and all DMTF information is available at **dmtf.org/standards/profiles/**.

SM-CLP commands implement a subset of the local RACADM commands. The commands are useful for scripting since you can execute these commands from a management station command line. You can retrieve the output of commands in well-defined formats, including XML, facilitating scripting and integration with existing reporting and management tools.

Topics:

- System management capabilities using SMCLP
- Running SMCLP commands
- iDRAC SMCLP syntax
- Navigating the map address space
- Using show verb
- Usage examples

System management capabilities using SMCLP

iDRAC SMCLP enables you to:

- Manage Server Power Turn on, shut down, or reboot the system
- Manage System Event Log (SEL) Display or clear the SEL records
- View iDRAC user accounts
- View system properties

Running SMCLP commands

You can run the SMCLP commands using SSH interface. Open an SSH and log in to iDRAC as an administrator. The SMCLP prompt (admin ->)is displayed.

SMCLP prompts:

- yx1x blade servers use -\$.
- yx1x rack and tower servers use admin->.
- yx2x blade, rack, and tower servers use admin->.



where, y is an alpha-numeric character such as M (for blade servers), R (for rack servers), and T (for tower servers) and x is a number. This indicates the generation of Dell PowerEdge servers.

(i) NOTE: Scripts using -\$ can use these for yx1x systems, but starting with yx2x systems one script with admin-> can be

used for blade, rack, and tower servers.

iDRAC SMCLP syntax

The iDRAC SMCLP uses the concept of verbs and targets to provide systems management capabilities through the CLI. The verb indicates the operation to perform, and the target determines the entity (or object) that runs the operation.

The SMCLP command line syntax:

<verb> [<options>] [<target>] [<properties>]

The following table provides the verbs and its definitions.

Table 62. SMCLP verbs

Verb	Definition
cd	Navigates through the MAP using the shell
set	Sets a property to a specific value
help	Displays help for a specific target
reset	Resets the target
show	Displays the target properties, verbs, and subtargets
start	Turns on a target
stop	Shuts down a target
exit	Exits from the SMCLP shell session
version	Displays the version attributes of a target
load	Moves a binary image to a specified target address from a URL

The following table provides a list of targets.

Table 63. SMCLP targets

Target	Definitions
admin1	admin domain
admin1/profiles1	Registered profiles in iDRAC
admin1/hdwr1	Hardware
admin1/system1	Managed system target
admin1/system1/capabilities1	Managed system SMASH collection capabilities



Table 63. SMCLP targets (continued)

Target	Definitions
admin1/system1/capabilities1/elecap1	Managed system target capabilities
admin1/system1/logs1	Record Log collections target
admin1/system1/logs1/log1	System Event Log (SEL) record entry
admin1/system1/logs1/log1/record*	An individual SEL record instance on the managed system
admin1/system1/settings1	Managed system SMASH collection settings
admin1/system1/capacities1	Managed system capacities SMASH collection
admin1/system1/consoles1	Managed system consoles SMASH collection
admin1/system1/sp1	Service Processor
admin1/system1/sp1/timesvc1	Service Processor time service
admin1/system1/sp1/capabilities1	Service processor capabilities SMASH collection
admin1/system1/sp1/capabilities1/clpcap1	CLP service capabilities
admin1/system1/sp1/capabilities1/ pwrmgtcap1	Power state management service capabilities on the system
admin1/system1/sp1/capabilities1/ acctmgtcap*	Account management service capabilities
admin1/system1/sp1/capabilities1/ rolemgtcap*	Local Role Based Management capabilities
admin1/system1/sp1/capabilities1/elecap1	Authentication capabilities
admin1/system1/sp1/settings1	Service Processor settings collection
admin1/system1/sp1/settings1/clpsetting1	CLP service settings data
admin1/system1/sp1/clpsvc1	CLP service protocol service
admin1/system1/sp1/clpsvc1/clpendpt*	CLP service protocol endpoint



Table 63. SMCLP targets (continued)

Target	Definitions
admin1/system1/sp1/clpsvc1/tcpendpt*	CLP service protocol TCP endpoint
admin1/system1/sp1/jobq1	CLP service protocol job queue
admin1/system1/sp1/jobq1/job*	CLP service protocol job
admin1/system1/sp1/pwrmgtsvc1	Power state management service
admin1/system1/sp1/account1-16	Local user account
admin1/sysetm1/sp1/account1-16/identity1	Local user identity account
admin1/sysetm1/sp1/account1-16/identity2	IPMI identity (LAN) account
admin1/sysetm1/sp1/account1-16/identity3	IPMI identity (Serial) account
admin1/sysetm1/sp1/account1-16/identity4	CLP identity account
admin1/system1/sp1/acctsvc2	IPMI account management service
admin1/system1/sp1/acctsvc3	CLP account management service
admin1/system1/sp1/rolesvc1	Local Role Base Authorization (RBA) service
admin1/system1/sp1/rolesvc1/Role1-16	Local role
admin1/system1/sp1/rolesvc1/Role1-16/ privilege1	Local role privilege
admin1/system1/sp1/rolesvc2	IPMI RBA service
admin1/system1/sp1/rolesvc2/Role1-3	IPMI role
admin1/system1/sp1/rolesvc2/Role4	IPMI Serial Over LAN (SOL) role
admin1/system1/sp1/rolesvc3	CLP RBA Service
admin1/system1/sp1/rolesvc3/Role1-3	CLP role



Table 63. SMCLP targets (continued)

Target	Definitions
admin1/system1/sp1/rolesvc3/Role1-3/ privilege1	CLP role privilege

Navigating the map address space

Objects that can be managed with SM-CLP are represented by targets arranged in a hierarchical space called the Manageability Access Point (MAP) address space. An address path specifies the path from the root of the address space to an object in the address space.

The root target is represented by a slash (/) or a backslash (\). It is the default starting point when you log in to iDRAC. Navigate down from the root using the cd verb.

NOTE: The slash (/) and backslash (\) are interchangeable in SM-CLP address paths. However, a backslash at the end of a command line continues the command on the next line and is ignored when the command is parsed.

For example to navigate to the third record in the System Event Log (SEL), enter the following command:

->cd /admin1/system1/logs1/log1/record3

Enter the cd verb with no target to find your current location in the address space. The .. and . abbreviations work as they do in Windows and Linux: .. refers to the parent level and . refers to the current level.

Using show verb

To learn more about a target use the show verb. This verb displays the target's properties, sub-targets, associations, and a list of the SM-CLP verbs that are allowed at that location.

Using the -display option

The show -display option allows you to limit the output of the command to one or more of properties, targets, associations, and verbs. For example, to display just the properties and targets at the current location, use the following command:

show -display properties, targets

To list only certain properties, qualify them, as in the following command:

show -d properties=(userid,name) /admin1/system1/sp1/account1

If you only want to show one property, you can omit the parentheses.

Using the -level option

The show -level option executes show over additional levels beneath the specified target. To see all targets and properties in the address space, use the -l all option.

Using the -output option

The -output option specifies one of four formats for the output of SM-CLP verbs: text, clpcsv, keyword, and clpxml.

The default format is **text**, and is the most readable output. The **clpcsv** format is a comma-separated values format suitable for loading into a spreadsheet program. The **keyword** format outputs information as a list of keyword=value pairs one per line. The **clpxml** format is an XML document containing a **response** XML element. The DMTF has specified the **clpcsv** and **clpxml** formats and their specifications can be found on the DMTF website at **dmtf.org**.

The following example shows how to output the contents of the SEL in XML:

show -1 all -output format=clpxml /admin1/system1/logs1/log1



Usage examples

This section provides use case scenarios for SMCLP:

- Server power management
- SEL management
- Map target navigation

Server power management

The following examples show how to use SMCLP to perform power management operations on a managed system.

Type the following commands at the SMCLP command prompt:

• To switch off the server:

stop /system1
The following message is displayed:
system1 has been stopped successfully

• To switch on the server:

start /system1

The following message is displayed:

system1 has been started successfully

• To reboot the server:

reset /system1

The following message is displayed:

system1 has been reset successfully

SEL management

The following examples show how to use the SMCLP to perform SEL-related operations on the managed system. Type the following commands at the SMCLP command prompt:

• To view the SEL:

show/system1/logs1/log1 The following output is displayed: /system1/logs1/log1 Targets: Record1 Record2 Record3 Record4 Record5 Properties: InstanceID = IPMI:BMC1 SEL Log MaxNumberOfRecords = 512 CurrentNumberOfRecords = 5 Name = IPMI SEL

324 Using SMCLP



```
EnabledState = 2
OperationalState = 2
HealthState = 2
Caption = IPMI SEL
Description = IPMI SEL
ElementName = IPMI SEL
Commands:
cd
show
help
exit
version
```

```
To view the SEL record:
•
  show/system1/logs1/log1
  The following output is displayed:
  /system1/logs1/log1/record4
  Properties:
  LogCreationClassName= CIM_RecordLog
  CreationClassName= CIM LogRecord
  LogName= IPMI SEL
  RecordID= 1
  MessageTimeStamp= 20050620100512.000000-000
  Description= FAN 7 RPM: fan sensor, detected a failure
  ElementName= IPMI SEL Record
  Commands:
  cd
  show
  help
  exit
  version
```

Map target navigation

The following examples show how to use the cd verb to navigate the MAP. In all examples, the initial default target is assumed to be /.

Type the following commands at the SMCLP command prompt:

• To navigate to the system target and reboot:

cd system1 reset The current default target is /.

• To navigate to the SEL target and display the log records:

```
cd system1
cd logs1/log1
```

show

• To display current target:

Using SMCLP 325



type cd .

- To move up one level: type cd ..
- To exit: exit



Deploying operating systems

You can use any of the following utilities to deploy operating systems to managed systems:

- Remote File Share
- Console

Topics:

- Deploying operating system using remote file share
- Deploying operating system using virtual media
- Deploying embedded operating system on SD card

Deploying operating system using remote file share

Before you deploy the operating system using Remote File Share (RFS), make sure that:

- Configure User and Access Virtual Media privileges for iDRAC are enabled for the user.
- Network share contains drivers and operating system bootable image file, in an industry standard format such as .img or .iso.

NOTE: While creating the image file, follow standard network-based installation procedures, and mark the deployment image as read-only to make sure that each target system boots and runs the same deployment procedure.

To deploy an operating system using RFS:

- 1. Using Remote File Share (RFS), mount the ISO or IMG image file to the managed system through NFS, CIFS, HTTP, or HTTPs.
 - () NOTE: RFS using HTTP, basic, or digest authentication is not supported, no authentication is requored. For HTTPS, basic authentication is not supported, only digest authentication or no authentication is supported.
- 2. Go to Configuration > System Settings > Hardware Settings > First Boot Device.
- 3. Set the boot order in the First Boot Device drop-down list to select a virtual media such as floppy, CD, DVD, or ISO.
- 4. Select the **Boot Once** option to enable the managed system to reboot using the image file for the next instance only.
- 5. Click Apply.
- 6. Reboot the managed system and follow the on-screen instructions to complete the deployment.

Managing remote file shares

Using Remote File Share (RFS) feature, you can set an ISO or IMG image file on a network share and make it available to the managed server's operating system as a virtual drive by mounting it as a CD or DVD using NFS, CIFS, HTTP or HTTPs. RFS is a licensed feature.

Remote file share supports only .img and .iso image file formats. A .img file is redirected as a virtual floppy and a .iso file is redirected as a virtual CDROM.

You must have Virtual Media privileges to perform an RFS mounting.

RFS and Virtual Media features are mutually exclusive.

- If the Virtual Media client is not active, and you attempt to establish an RFS connection, the connection is established and the remote image is available to the host operating system.
- If the Virtual Media client is active, and you attempt to establish an RFS connection, the following error message is displayed:

Virtual Media is detached or redirected for the selected virtual drive.



The connection status for RFS is available in iDRAC log. Once connected, an RFS-mounted virtual drive does not disconnect even if you log out from iDRAC. The RFS connection is closed if iDRAC is reset or the network connection is dropped. The Web interface and command-line options are also available in CMCOME Modular and iDRAC to close the RFS connection. The RFS connection from CMC always overrides an existing RFS mount in iDRAC.

() NOTE:

- CIFS and NFS supports both IPv4 and IPv6 addresses.
- When the iDRAC is configured with both IPv4 and IPv6, the DNS server can contain records associating the iDRAC hostname to both addresses. If IPv4 option is disabled in iDRAC, then iDRAC may not be able to access the external IPv6 share. This is because the DNS server may still contain IPv4 records, and DNS name resolution can return the IPv4 address. In such cases, it is recommended to delete the IPv4 DNS records from the DNS server, when disabling IPv4 option in iDRAC.
- If you are using CIFS and are part of an Active Directory domain, enter the domain name with the IP address in the image file path.
- If you want to access a file from an NFS share, configure the following share permissions. These permissions are required because iDRAC interfaces run in non-root mode.
 - Linux: Ensure that the share permissions are set to at least **Read** for the **Others** account.
 - Windows: Go to the **Security** tab of the share properties and add **Everyone** to **Groups or user names** field with **Read & execute** privilege.
- If ESXi is running on the managed system and if you mount a floppy image (**.img**) using RFS, the connected floppy image is not available to the ESXi operating system.
- iDRAC vFlash feature and RFS are not related.
- Only English ASCII characters are supported in network share file paths.
- The OS drive eject feature is not supported when virtual media is connected using RFS.
- RFS through HTTP or HTTPs feature is not available on CMC web interface.

Configuring remote file share using web interface

To enable remote file sharing:

- 1. In iDRAC web interface, go to **Configuration** > **Virtual Media** > **Attached Media**. The **Attached Media** page is displayed.
- 2. Under Attached Media, select Attach or Auto Attach.
- 3. Under **Remote File Share**, specify the image file path, domain name, user name, and password. For information about the fields, see the *iDRAC Online Help*.

Example for image file path:

- $\bullet~\mbox{CIFS} //<\mbox{IP}$ to connect for CIFS file system>/<file path>/<image name>
- NFS < IP to connect for NFS file system>:/<file path>/<image name>
- HTTP http://<URL>/<file path>/<image name>
- HTTPs https://<URL>/<file path>/<image name>

NOTE: To avoid I/O errors when using CIFS shares hosted on Windows 7 systems, modify the following registry keys:

- Set HKLM\SYSTEM\CurrentControlSet\Control\Session Manager\Memory Management\LargeSystemCache to 1
- Set HKLM\SYSTEM\CurrentControlSet\Services\LanmanServer\Parameters\Size to 3

i NOTE: Both '/' or '\' characters can be used for the file path.

CIFS supports both IPv4 and IPv6 addresses but NFS supports only IPv4 address.

If you are using NFS share, make sure that you provide the exact <file path> and <image name> as it is case-sensitive.

NOTE: For information on recommended characters for user names and passwords, see Recommended characters in user names and passwords.



() NOTE: The characters allowed in user names and passwords for network shares are determined by the network-share type. iDRAC supports valid characters for network share credentials as defined by the share type, except <, >, and , (comma).

4. Click Apply and then click Connect.

After the connection is established, the Connection Status displays Connected.

NOTE: Even if you have configured remote file sharing, the Web interface does not display user credential information due to security reasons.

NOTE: If the image path contains user credentials, use HTTPS to avoid credentials from displaying in the GUI and RACADM. If entering the credentials in the URL, avoid using "@" symbol, because it is a separator character.

For Linux distributions, this feature may require a manual mount command when operating at runlevel init 3. The syntax for the command is:

mount /dev/OS_specific_device / user_defined_mount_point

Where, user defined mount point is any directory you choose to use for the mount similar to any mount command.

For RHEL, the CD device (.iso virtual device) is /dev/scd0 and floppy device (.img virtual device) is /dev/sdc.

For SLES, the CD device is /dev/sr0 and the floppy device is /dev/sdc. To make sure that the correct device is used (for either SLES or RHEL), when you connect the virtual device, on the Linux OS you must immediately run the command:

tail /var/log/messages | grep SCSI

This displays the text that identifies the device (example, SCSI device sdc). This procedure also applies to Virtual Media when you are using Linux distributions in runlevel init 3. By default, the virtual media is not auto-mounted in init 3.

Configuring remote file share using RACADM

To configure remote file share using RACADM, use:

racadm remoteimage

```
racadm remoteimage <options>
```

Options are:

- -c: connect image
- -d : disconnect image
- -u <username>: username to access the network share
- -p <password>: password to access the network share

-1 <image_location>: image location on the network share; use double quotes around the location. See examples for image file path in Configuring Remote File Share Using Web Interface section

- -s : display current status
- () NOTE: All characters including alphanumeric and special characters are allowed as part of user name, password, and image_location except the following characters: ' (single quote), "(double quote), ,(comma), < (less than), and > (greater than).

() NOTE: To avoid I/O errors when using CIFS shares hosted on Windows 7 systems, modify the following registry keys:

- Set HKLM\SYSTEM\CurrentControlSet\Control\Session Manager\Memory Management\LargeSystemCache to 1
- Set HKLM\SYSTEM\CurrentControlSet\Services\LanmanServer\Parameters\Size to 3

Deploying operating system using virtual media

Before you deploy the operating system using Virtual Media, make sure that:

Virtual Media is in Attached state for the virtual drives to appear in the boot sequence.



- If Virtual Media is in Auto Attached mode, the Virtual Media application must be launched before booting the system.
- Network share contains drivers and operating system bootable image file, in an industry standard format such as .img or .iso.

To deploy an operating system using Virtual Media:

- **1.** Do one of the following:
 - Insert the operating system installation CD or DVD into the management station CD or DVD drive.
 - Attach the operating system image.
- 2. Select the drive on the management station with the required image to map it.
- 3. Use one of the following methods to boot to the required device:
 - Set the boot order to boot once from Virtual Floppy or Virtual CD/DVD/ISO using the iDRAC Web interface.
 - Set the boot order through **System Setup** > **System BIOS Settings** by pressing <F2> during boot.
- 4. Reboot the managed system and follow the on-screen instructions to complete the deployment.

Installing operating system from multiple disks

- 1. Unmap the existing CD/DVD.
- 2. Insert the next CD/DVD into the remote optical drive.
- **3.** Remap the CD/DVD drive.

Deploying embedded operating system on SD card

To install an embedded hypervisor on an SD card:

- 1. Insert the two SD cards in the Internal Dual SD Module (IDSDM) slots on the system.
- 2. Enable SD module and redundancy (if required) in BIOS.
- 3. Verify if the SD card is available on one of the drives when you <F11> during boot.
- 4. Deploy the embedded operating system and follow the operating system installation instructions.

Enabling SD module and redundancy in BIOS

To enable SD module and redundancy in BIOS:

- **1.** Press <F2> during boot.
- 2. Go to System Setup > System BIOS Settings > Integrated Devices.
- 3. Set the Internal USB Port to On. If it is set to Off, the IDSDM is not available as a boot device.
- If redundancy is not required (single SD card), set Internal SD Card Port to On and Internal SD Card Redundancy to Disabled.
- 5. If redundancy is required (two SD cards), set Internal SD Card Port to On and Internal SD Card Redundancy to Mirror.
- 6. Click **Back** and click **Finish**.
- 7. Click Yes to save the settings and press <Esc> to exit System Setup.

About IDSDM

Internal Dual SD Module (IDSDM) is available only on applicable platforms. IDSDM provides redundancy on the hypervisor SD card by using another SD card that mirrors the first SD card's content.

Either of the two SD cards can be the master. For example, if two new SD cards are installed in the IDSDM, SD1 is active (master) card and SD2 is the standby card. The data is written on both the cards, but the data is read from SD1. At any time if SD1 fails or is removed, SD2 automatically become the active (master) card.

You can view the status, health, and the availability of IDSDM using iDRAC Web Interface or RACADM. The SD card redundancy status and failure events are logged to SEL, displayed on the front panel, and PET alerts are generated if alerts are enabled.



Troubleshooting managed system using iDRAC

You can diagnose and troubleshoot a remote managed system using:

- Diagnostic console
- Post code
- Boot and crash capture videos
- Last system crash screen
- System event logs
- Lifecycle logs
- Front panel status
- Trouble indicators
- System health

Topics:

- Using diagnostic console
- Viewing post codes
- Viewing boot and crash capture videos
- Viewing logs
- Viewing last system crash screen
- Viewing System status
- Hardware trouble indicators
- Viewing system health
- Checking server status screen for error messages
- Restarting iDRAC
- Reset to Custom Defaults (RTD)
- Erasing system and user data
- Resetting iDRAC to factory default settings

Using diagnostic console

iDRAC provides a standard set of network diagnostic tools that are similar to the tools included with Microsoft Windows or Linux-based systems. Using iDRAC Web interface, you can access the network debugging tools.

To access Diagnostics Console:

- 1. In the iDRAC Web interface, go to Maintenance > Diagnostics. The Diagnostics Console Command page is displayed.
- 2. In the **Command** text box, enter a command and click **Submit**. For information about the commands, see the *iDRAC Online Help*.

The results are displayed on the same page.

Reset iDRAC and Reset iDRAC to default

- 1. In the iDRAC Web interface, go to **Maintenance** > **Diagnostics**. You have the following options:
 - Click **Reset iDRAC** to reset the iDRAC. A normal reboot operation is performed on the iDRAC. After reboot, refresh the browser to reconnect and log in to iDRAC.

Troubleshooting managed system using iDRAC 331



- Click **Reset iDRAC to Default Settings** to reset the iDRAC to the default settings. After you click **Reset iDRAC to Default Settings**, **Reset iDRAC to factory default** window is displayed. This action reset the iDRAC to the factory defaults. Chose any of the following options:
 - **a.** Preserve user and network settings.
 - b. Discard all settings and reset users to the shipping value (root/shipping value).
- c. Discard all settings and reset username and password.
- 2. A warning message is displayed. Click \mathbf{Ok} to proceed further.

Scheduling remote automated diagnostics

You can remotely invoke automated offline diagnostics on a server as a one-time event and return the results. If the diagnostics require a reboot, you can reboot immediately or stage it for a subsequent reboot or maintenance cycle (similar to updates). When diagnostics are run, the results are collected and stored in the internal iDRAC storage. You can then export the results to an NFS, CIFS, HTTP, or HTTPs network share using the diagnostics export racadm command. You can also run diagnostics using the appropriate WSMan command(s). For more information, see the WSMan documentation.

You must have iDRAC Express license to use remote automated diagnostics.

You can perform the diagnostics immediately or schedule it on a particular day and time, specify the type of diagnostics, and the type of reboot.

For the schedule, you can specify the following:

- Start time Run the diagnostic at a future day and time. If you specify TIME NOW, the diagnostic is run on the next reboot.
- End time Run the diagnostic until a date and time after the Start time. If it is not started by End time, it is marked as failed with End time expired. If you specify TIME NA, then the wait time is not applicable.

The types of diagnostic tests are:

- Express test
- Extended test
- Both in a sequence

The types of reboot are:

- Power cycle system
- Graceful shutdown (waits for operating system to turn off or for system restart)
- Forced Graceful shutdown (signals operating system to turn off and waits for 10 minutes. If the operating system does not turn off, the iDRAC power cycles the system)

Only one diagnostic job can be scheduled or run at one time. A diagnostic job can complete successfully, complete with error, or is unsuccessful. The diagnostic events including the results are recorded in Lifecycle Controller log. You can retrieve the results of the last diagnostic execution using remote RACADM or WSMan.

You can export the diagnostic results of the last completed diagnostics that were scheduled remotely to a network share such as CIFS, NFS, HTTP or HTTPS. The maximum file size is 5 MB.

You can cancel a diagnostic job when the status of the job is Unscheduled or Scheduled. If the diagnostic is running, then restart the system to cancel the job.

Before you run the remote diagnostics, make sure that:

- Lifecycle Controller is enabled.
- You have Login and Server Control privileges.

Scheduling remote automated diagnostics using RACADM

• To run the remote diagnostics and save the results on the local system, use the following command:

racadm diagnostics run -m <Mode> -r <reboot type> -s <Start Time> -e <Expiration Time>

• To export the last run remote diagnostics results, use the following command:

```
racadm diagnostics export -f <file name> -l <NFS / CIFS / HTTP / HTTPs share> -u
<username> -p <password>
```

For more information about the options, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.



Viewing post codes

Post codes are progress indicators from the system BIOS, indicating various stages of the boot sequence from power-on-reset, and allows you to diagnose any faults related to system boot-up. The **Post Codes** page displays the last system post code prior to booting the operating system.

To view the Post Codes, go to Maintenance > Troubleshooting > Post Code.

The Post Code page displays the system health indicator, a hexadecimal code, and a description of the code.

Viewing boot and crash capture videos

You can view the video recordings of:

- Last three boot cycles A boot cycle video logs the sequence of events for a boot cycle. The boot cycle videos are arranged in the order of latest to oldest.
- Last crash video A crash video logs the sequence of events leading to the failure.

This is a licensed feature.

iDRAC records fifty frames during boot time. Playback of the boot screens occur at a rate of 1 frame per second. If iDRAC is reset, the boot capture video is not available as it is stored in RAM and is deleted.

() NOTE:

- You must have Access Virtual Console or administrator privileges to playback the Boot Capture and Crash Capture videos.
- The video capture time displayed in the iDRAC GUI video player may differ from the video capture time displayed in other video players. The iDRAC GUI video player displays the time in the iDRAC time zone while all other video players display the time in the respective operating system time zones.
- () NOTE: DVC boot capture files are not videos. They are sequence of screens (at 1 particular resolution) taken during the course of the server boot. The DVC player converts these screens together to create the boot video. When you export the video from DVC (continuous snapshot and differences) to .mov (actual video) format, it is expected to use the same resolution, or a similar resolution, that the video was initially encoded with. Videos need to be exported at a similar resolution that they have been captured with.
- **NOTE:** The reason for the delay in boot capture file availability is because the boot capture buffer is not full after the host boot.

To view the Boot Capture screen, click Maintenance > Troubleshooting > Video Capture.

The Video Capture screen displays the video recordings. For more information, see the iDRAC Online Help.

NOTE: When embedded video controller is disabled and server has add-on video controller, then certain latency is expected with respect to boot capture. Hence, End of Post Messages of a video will be recorded in next capture.

Configuring video capture settings

To configure the video capture settings:

- In the iDRAC Web interface, go to Maintenance > Troubleshooting > Video Capture. The Video Capture page is displayed.
- 2. From the Video Capture Settings drop-down menu, select any of the following options:
 - **Disable** Boot capture is disabled.
 - Capture until buffer full Boot sequence is captured until the buffer size has reached.
 - Capture until end of POST Boot sequence is captured until end of POST.
- 3. Click Apply to apply the settings.



Viewing logs

You can view System Event Logs (SELs) and Lifecycle logs. For more information, see Viewing System Event Log and Viewing Lifecycle log.

Viewing last system crash screen

The last crash screen feature captures a screenshot of the most recent system crash, saves, and displays it in iDRAC. This is a licensed feature.

To view the last crash screen:

- 1. Make sure that the last system crash screen feature is enabled.
- 2. In iDRAC Web interface, go to Overview > Server > Troubleshooting > Last Crash Screen.

The Last Crash Screen page displays the last saved crash screen from the managed system.

Click **Clear** to delete the last crash screen.

i NOTE: Once iDRAC is reset or an AC power cycle event occurs, then the crash capture data is cleared.

Viewing System status

The System Status summarizes the status of the following components in the system:

- Summary
- Batteries
- Cooling
- CPUs
- Front Panel
- Intrusion
- Memory
- Network Devices
- Power Supplies
- Voltages
- Removable Flash Media
- Chassis Controller

You can view the status of the managed system:

- For rack and tower servers: LCD front panel and system ID LED status or LED front panel and system ID LED status.
- For blade servers: Only system ID LEDs.

Viewing system front panel LCD status

To view the LCD front panel status for applicable rack and tower servers, in iDRAC Web interface, go to **System** > **Overview** > **Front Panel**. The **Front Panel** page is displayed.

The **Front Panel** section displays the live feed of the messages currently being displayed on the LCD front panel. When the system is operating normally (indicated by solid blue color in the LCD front panel), both **Hide Error** and **UnHide Error** are grayed-out.

(i) NOTE: You can hide or unhide the errors only for rack and tower servers.

Based on the selection, the text box displays the current value. If you select User Defined, enter the required message in the text box. The character limit is 62. If you select None, home message is not displayed on the LCD.

To view LCD front panel status using RACADM, use the objects in the System.LCD group. For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.



Viewing system front panel LED status

To view the current system ID LED status, in iDRAC web interface, go to **System** > **Overview** > **Front Panel**. The **Front Panel** section displays the current front panel status:

- Solid blue No errors present on the managed system.
- Blinking blue Identify mode is enabled (regardless of managed system error presence).
- Solid amber Managed system is in failsafe mode.
- Blinking amber Errors present on managed system.

When the system is operating normally (indicated by blue Health icon on the LED front panel), then both **Hide Error** and **UnHide Error** is grayed-out. You can hide or unhide the errors only for rack and tower servers.

To view system ID LED status using RACADM, use the getled command.

For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Hardware trouble indicators

The hardware related problems are:

- Failure to power up
- Noisy fans
- Loss of network connectivity
- Hard drive failure
- USB media failure
- Physical damage

Based on the problem, use the following methods to correct the problem:

- Reseat the module or component and restart the system
- In case of a blade server, insert the module into a different bay in the chassis
- Replace hard drives or USB flash drives
- Reconnect or replace the power and network cables

If problem persists, see the *Installation and Service Manual* available at https://www.dell.com/poweredgemanuals for specific troubleshooting information about the hardware device.

CAUTION: You should only perform troubleshooting and simple repairs as authorized in your product documentation, or as directed by online or telephone service and support team. Damage due to servicing that is not authorized by Dell is not covered by your warranty. Read and follow the safety instructions that came with the product.

Viewing system health

You can view the status for the following components on the iDRAC, CMC and, OME-Modular Web interfaces:

- Batteries
- CPUs
- Cooling
- Intrusion
- Memory
- Power Supplies
- Removable Flash Media
- Voltages
- Miscellaneous

Click any component name in the **Server Health** section to view details about the component.



Checking server status screen for error messages

When a flashing amber LED is blinking, and a particular server has an error, the main Server Status Screen on the LCD highlights the affected server in orange. Use the LCD navigation buttons to highlight the affected server, then click the center button. Error and warning messages will be displayed on the second line. For the list of error messages displayed on the LCD panel, see the server's Owner's Manual.

Restarting iDRAC

You can perform a hard or soft iDRAC restart without turning off the server:

- Hard restart On the server, press and hold the LED button for 15 seconds.
- Soft restart Using iDRAC Web interface or RACADM.

Reset to Custom Defaults (RTD)

You can use Reset to Custom Defaults feature to upload a custom config file and RTD to the settings. The new settings are applied on top of preserving users and network settings.

Reset to Custom Defaults feature has following options:

- Upload Custom Default Settings
 - You can upload custom defaults settings file. This file can be obtained by exporting the Server Configuration Profile(SCP) in XML format (JSON format is not supported for this feature). The contents of the file can be modified by customer to add or delete the settings.
 - \circ $\,$ You can upload the SCP XML file using iDRAC GUI or RACADM interfaces.
 - The uploaded configurations are saved in the default database.
- Save current settings as custom defaults
 - This operation saves the current settings as default settings.
 - This is only supported via RACADM interface.
- Download custom default settings
 - You can download SCP XML for all the default settings.
 - This is only supported via RACADM interface.
- Initiate reset to custom defaults —
- The uploaded/saved default settings are applied.

Resetting iDRAC using iDRAC web interface

To reset iDRAC, do one of the following in the iDRAC Web interface:

- Upload Custom Defaults file:
 - Go to Configuration > Server Configuration Profile > Custom Defaults > Upload Custom Defaults
 - Upload the customized CustomConfigured.xml file from Local Share path
 - Click **Apply**. New Upload Custom Defaults Job is created.
- Reset to Custom Defaults:
 - When Upload Custom Defaults job is successful, go to Maintainance > Diagnostics, click Reset iDRAC to Factory Defalts option.
 - $\circ~$ Select Discard all settings and set to Custom default configuration.
 - Click **Continue** to initiate Reset to customs Defaults configuration.

Resetting iDRAC using RACADM

To restart iDRAC, use the **racreset** command. For more information, see the *Chassis Management Controller RACADM CLI Guide* available at https://www.dell.com/cmcmanuals.For more information, see the *OME* - *Modular for PowerEdge MX7000 Chassis RACADM CLI Guide* available at https://www.dell.com/openmanagemanuals

For Reset to default operations, use the following commands:



- Upload Custom Defaults file racadm -r <iDracIP> -u <username> -p <Password> set -f <filename> -t xml --customdefaults
- Save Current Settings as Default settings racadm -r <iDracIP> -u <username> -p <Password> set --savecustomdefaults
- Download Custom Default settings racadm -r <iDracIP> -u <username> -p <Password> get -f <filename> -t xml --customdefaults
- Reset to Custom Defaults Racadm -r <iDracIP> -u <username> -p <Password> racresetcfg -custom

Erasing system and user data

(i) NOTE: Erasing system and user data is not supported from iDRAC GUI.

You can erase system component(s) and user data for the following components:

- BIOS reset to default
- Embedded Diagnostics
- Embedded OS Driver Pack
- Lifecycle Controller Data
- iDRAC reset to default
- Overwrite hard drives that do not support Instant Secure Erase (ISE)
- Reset controller cache
- Reset vFLASH
- Erase Hard Drives, SSDs, and NVMes that support ISE
- Clear all OS applications

Before performing system erase, ensure that:

You have iDRAC Server Control privilege.

• Lifecycle Controller is enabled.

The Lifecycle Controller Data option erases any content such as the LC Log, configuration database, rollback firmware, factory as-shipped logs, and the configuration information from the FP SPI (or management riser).

NOTE: The Lifecycle Controller log contains the information about the system erase request and any information generated when the iDRAC restarts. All previous information is removed.

You can delete individual or multiple system components using the SystemErase command:

racadm systemErase <BIOS | DIAG | DRVPACK | LCDATA | IDRAC >

where,

- bios BIOS reset to default
- diag Embedded Diagnostics
- drvpack Embedded OS Driver Pack
- Icdata Clear the Lifecycle Controller Data
- idrac iDRAC reset to default
- overwritepd Overwrite hard drives that do not support Instant Secure Erase (ISE)
- percnvcache Reset controller cache
- vflash Reset vFLASH
- secureerasepd Erase Hard Drives, SSDs, and NVMes that support ISE
- allapps Clears all OS applications

(i) NOTE: While erasing vFlash, ensure that all partitions on the vFlash card are detached before executing the operation.

() NOTE: If SEKM is enabled on the server, then disable SEKM using the racadm sekm disable command before using this command. This can avoid any storage devices being locked out which are secured by iDRAC, if SEKM settings are erased from iDRAC by executing this command.

For more information, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/idracmanuals.

Troubleshooting managed system using iDRAC 337



- **NOTE:** The Dell tech center link appears on the iDRAC GUI on Dell branded systems. If you erase system data by using WSMan command and want the link to appear again, reboot the host manually and wait for CSIOR to run.
- **NOTE:** After you run System Erase, the VDs may still appear. Run CSIOR after System Erase is completed and iDRAC is rebooted.

Resetting iDRAC to factory default settings

You can reset iDRAC to the factory default settings using the iDRAC Settings utility or the iDRAC Web interface.

Resetting iDRAC to factory default settings using iDRAC web interface

To reset iDRAC to factory default settings using the iDRAC Web interface:

- 1. Go to Maintenance > Diagnostics. The Diagnostics Console page is displayed.
- Click Reset iDRAC to Default Settings. The completion status is displayed in percentage. iDRAC reboots and is restored to factory defaults. The iDRAC IP is reset and is not accessible. You can configure the IP using the front panel or BIOS.

Resetting iDRAC to factory default settings using iDRAC settings utility

To reset iDRAC to factory default values using the iDRAC Settings utility:

- Go to Reset iDRAC configurations to defaults. The iDRAC Settings Reset iDRAC configurations to defaults page is displayed.
- 2. Click Yes.
- iDRAC reset starts.
- 3. Click **Back** and navigate to the same **Reset iDRAC configurations to defaults** page to view the success message.



SupportAssist Integration in iDRAC

SupportAssist allows you to create SupportAssist collections and utilize other SupportAssist features to monitor your system and datacenter. iDRAC provides an application interfaces for gathering platform information that enables support services to resolve platform and system problems. iDRAC helps you to generate a SupportAssist collection of the server and then export the collection to a location on the management station (local) or to a shared network location such as FTP, Trivial File Transfer Protocol (TFTP), HTTP, HTTPS, Common Internet File System (CIFS) or Network File Share (NFS). The collection is generated in the standard ZIP format. You can send this collection to technical support for troubleshooting or inventory collection. **Topics:**

- SupportAssist Registration
- Installing Service Module
- Server OS Proxy Information
- SupportAssist
- Service Request Portal
- Collection Log
- Generating SupportAssist Collection
- Settings
- Collection Settings
- Contact Information

SupportAssist Registration

To take advantage of the automated, proactive, and predictive features of SupportAssist, you must register your system with SupportAssist.

You can generate and save a collection locally or to a network, and also send to Dell EMC without registration.

NOTE: Some OEM customers do not have the Model name. The back-end Support Assist does not allow registering such systems with DELL.

Contact and shipping information

To complete the registration, you must provide the contact and shipping information.

Primary contact information

Enter Company name, Country, First Name*, Last Name*, Phone Number*, Alternate Number, and Email Address*. Check if the details are displayed correctly and make changes if you want to edit any of the fields.

*indicates that the fields are mandatory.

Secondary contact information

Enter First Name, Last name, Phone Number, Alternate Number, Email Address and check if the details are displayed correctly and make changes if you want to edit any of the fields.



SupportAssist Integration in iDRAC 339



Auto dispatch

When a critical event is reported to Dell-EMC through iDRAC that is registered for SupportAssist, auto dispatch workflow may be initiated. This workflow is based on the event being forwarded and registered device SupportAssist warranty level. You must enter the **Dispatch information** during the SupportAssist registration process to enable auto dispatch workflow. If onsite support is required along with dispatch parts then select **Parts Dispatch with Onsite Support**.

NOTE: Auto dispatch is enabled in systems with iDRAC Service Module (iSM) v3.4.0 for Windows. Future iSM releases will support auto dispatch for additional operating systems.

Dispatch Address

Enter an address and the preferred contact hours.

End-user license agreement

After providing all the required information, you need to accept the End User License Agreement (EULA) to complete the registration process. You have the option to print the EULA for further references. You can cancel and terminate the registration process at any point of time.

Installing Service Module

In order to register and use SupportAssist, you must have iDRAC Service Module (iSM) installed in the system. Once you **initiate Service Module Installation** you can see the installation instructions. The **Next** button remains disabled until you successfully install iSM.

Server OS Proxy Information

In case there is an issue with the connection, then the user will be prompted to provide OS proxy information. Enter **Server**, **Port**, **Username** and **Password** to configure the proxy settings.

SupportAssist

Once SupportAssist is configured, you can check the SupportAssist dash board to view the **Service Request Summary**, **Warranty Status**, **SupportAssist Overview**, **Service Requests**, and **Collection log**. Registration is not required to view or send the Collection log.

Service Request Portal

Service Request shows the Status (Open/Closed), Description, Source (Event/Phone), Service Request ID, Date Opened and Date Closed details for each event. You can select and view further details of each event. You have the option to check Service Request Portal to view additional information for any individual case.

Collection Log

Collection Log shows the details of Collection Date and Time, Collection Type (Manual, Scheduled, Event based), Data Collected (Custom Selection, All Data), Collection Status (Complete with Errors, Complete), Job ID, Sent Status, and Sent Date and Time. You can send the last persisted collection in iDRAC to Dell.

NOTE: Once generated, the Collection Log Details can be filtered to remove the Personally Identifiable Information (PII) based on the user selection.

340 SupportAssist Integration in iDRAC

Inserido ao protocolo **22.951.206-4** por: **Pedro Henrique Golin Linhares** em: 31/03/2025 12:49. A autenticidade deste documento pode ser validada no endereço: https://www.eprotocolo.pr.gov.br/spiweb/validarDocumento com o código: fc1c44dfe0686d5547ffd691199be699.



Generating SupportAssist Collection

For generating the OS and Application logs:

- iDRAC Service Module must be installed and running in Host Operating System.
- OS Collector, which comes factory installed in iDRAC, if removed must be installed in iDRAC.

If you have to work with Tech Support on an issue with a server but the security policies restrict direct internet connection, then you can provide Tech Support with necessary data to facilitate troubleshooting of the problem without having to install software or download tools from Dell and without having access to the Internet from the server operating system or iDRAC.

You can generate a health report of the server and then export the Collection log:

- To a location on the management station (local).
- To a shared network location such as Common Internet File System (CIFS) or Network File Share (NFS). To export to a
- network share such as CIFS or NFS, direct network connectivity to the iDRAC shared or dedicated network port is required.
- To Dell EMC.

The SupportAssist Collection is generated in the standard ZIP format. The collection may contain the following information:

- Hardware inventory for all components (includes system component configuration and firmware details, Motherboard System Event Logs, iDRAC state information and Lifecycle Controller logs).
- Operating system and application information.
- Storage Controller logs.
- iDRAC Debug Logs.
- It contains an HTML5 viewer, that can be accessed once the collection is complete.
- The collection provides a massive amount of detailed system information and logs in a user friendly format that can be viewed without uploading the collection to the Tech Support site.

After the data is generated, you can view the data which contains multiple XML files and log files.

Each time the data collection is performed, an event is recorded in the Lifecycle Controller log. The event includes information such as the user who initiated the report, interface used, and the date and time of export.

On Windows, If WMI is disabled, OS Collector collection stops with an error message.

Check the appropriate privilege levels and make sure there is no firewall or security settings that may prevent from collecting the registry or software data.

Before generating the health report, make sure:

- Lifecycle Controller is enabled.
- Collect System Inventory On Reboot (CSIOR) is enabled.
- You have Login and Server Control privileges.

Generating SupportAssist Collection manually using iDRAC web interface

To generate the SupportAssist collection manually:

- 1. In the iDRAC Web interface, go to Maintenance > SupportAssist.
- 2. If the server is not registered for SupportAssist, SupportAssist Registration Wizard is displayed. Click Cancel > Cancel Registration.
- 3. Click Start a Collection.
- 4. Select the data sets to be included in the Collection.
- 5. You can opt to filter the collection for PII.
- 6. Select the destination where Collection needs to be saved.
 - a. If the server is connected to the internet and the **Send Now** option is enabled, then selecting this option transmits the Collection log to Dell EMC SupportAssist.
 - b. Save locally option allows you to save the generated Collection in the local system.
 - c. Save to Network option saves the generated Collection to user defined CIFS or NFS share location.
 i) NOTE: If Save to Network is selected, and no default location is available, the provided network details will be saved as default location for future collections. If default location already exist, then the collection will use the details specified once only.



If **Save to Network** option is selected, the user provided network details is saved as defaults (if no prior network share location have been saved) for any future collections.

- 7. Click **Collect** to proceed with Collection generation.
- 8. If prompted, accept the End User Level Agreement (EULA) to continue.
 - OS and Application Data option is grayed out and not selectable if:
 - iSM is not installed or running in Host OS, or
 - OS Collector has been removed from iDRAC, or
 - OS-BMC pass through is disabled in iDRAC, or
 - cached OS Application data is not available in iDRAC from a previous collection

Settings

This page allows you to configure the collection log settings, and if registered, you can update the contact details, enable or disable email notifications, and change the language settings.

Collection Settings

You can save the collections to a preferred network location. Use **Set Archive Directory** to set the network location. You can save the collections to a preferred network location. Use Set Archive Directory to set the network location. Enter the type of Protocol (CIFS/NFS) that you want to opt for, corresponding IP Address, Share Name, Domain Name, User Name and Password before you Test Network Connection. The Test Network Connection button will confirm a connection to the destination share.

If registered, you can opt to include Identification information while sending the data to Dell in the Collection Settings.

You can enable and schedule **Automatic Collection** options to avoid any manual intervention and keep a periodical check of the system. By default, when an event is triggered and support case is opened, SupportAssist is configured to automatically collect the system logs from the device that generated the alert and upload it to Dell. You can enable or disable Automatic Collection based on events. You can schedule the Automatic collections based on your suitable requirements. The available options are, weekly, monthly, quarterly or never. You can also configure the date and time for the scheduled periodic events. You have the option to enable or disable **ProSupport Plus Recommendation Report** while configuring the Automatic Collections.

Contact Information

This page shows the contact information details that were added during the registration of SupportAssist, and allows you to update them.



Frequently asked questions

This section lists the frequently asked questions for the following:

- System Event Log
- Network security
- Active Directory
- Single Sign On
- Smart card login
- Virtual console
- Virtual media
- vFlash SD card
- SNMP authentication
- Storage devices
- iDRAC Service Module
- RACADM
- Miscellaneous

Topics:

- System Event Log
- Custom sender email configuration for iDRAC alerts
- Network security
- Telemetry streaming
- Active Directory
- Single Sign-On
- Smart card login
- Virtual console
- Virtual media
- vFlash SD card
- SNMP authentication
- Storage devices
- GPU (Accelerators)
- iDRAC Service Module
- RACADM
- Permanently setting the default password to calvin
- Miscellaneous

System Event Log

While using iDRAC Web interface through Internet Explorer, why does SEL not save using the Save As option?

This is due to a browser setting. To resolve this:

1. In Internet Explorer, go to **Tools** > Internet Options > Security and select the zone you are attempting to download in.

For example, if the iDRAC device is on the local intranet, select **Local Intranet** and click **Custom level...**.

- 2. In the Security Settings window, under Downloads make sure that the following options are enabled:
 - Automatic prompting for file downloads (if this option is available)
 - File download

CAUTION: To make sure that the computer used to access iDRAC is safe, under Miscellaneous, do not enable the Launching applications and unsafe files option.



Custom sender email configuration for iDRAC alerts

Alert generated email is not from Custom sender email set on Cloud based email service.

You need to register your cloud email through this process : Support.google.com.

Network security

While accessing the iDRAC Web interface, a security warning is displayed stating that the SSL certificate issued by the Certificate Authority (CA) is not trusted.

iDRAC includes a default iDRAC server certificate to ensure network security while accessing through the Web-based interface and remote RACADM. This certificate is not issued by a trusted CA. To resolve this, upload a iDRAC server certificate issued by a trusted CA (for example, Microsoft Certificate Authority, Thawte or Verisign).

Why the DNS server not registering iDRAC?

Some DNS servers register iDRAC names that contain only up to 31 characters.

When accessing the iDRAC Web-based interface, a security warning is displayed stating that the SSL certificate hostname does not match the iDRAC hostname.

iDRAC includes a default iDRAC server certificate to ensure network security while accessing through the Web-based interface and remote RACADM. When this certificate is used, the web browser displays a security warning because the default certificate that is issued to iDRAC does not match the iDRAC hostname (for example, the IP address).

To resolve this, upload an iDRAC server certificate issued to the IP address or the iDRAC hostname. When generating the CSR (used for issuing the certificate), ensure that the common name (CN) of the CSR matches the iDRAC IP address (if certificate issued to IP) or the registered DNS iDRAC name (if certificate is issued to iDRAC registered name).

To make sure that the CSR matches the registered DNS iDRAC name:

1. In iDRAC Web interface, go to **Overview** > **iDRAC Settings** > **Network**. The **Network** page is displayed.

- 2. In the Common Settings section:
 - Select the **Register iDRAC on DNS** option.
 - In the DNS iDRAC Name field, enter the iDRAC name.
- 3. Click Apply.

Why am I unable to access iDRAC from my web browser?

This issue may occur if HTTP Strict Transport Security (HSTS) is enabled. HSTS is a web security mechanism which allows web browsers to interact using only the secure HTTPS protocol, and not HTTP.

Enable HTTPS on your browser and login to iDRAC to resolve the issue.

Why am I unable to complete operations that involve a remote CIFS share?

Import/export or any other remote file share operations that involve a CIFS share fail if they use only SMBv1. Ensure that the SMBv2 protocol is enabled on the server providing SMB/CIFS share. Refer to the Operating System documentation on how to enable the SMBv2 protocol.

Telemetry streaming

Few report data are missing while streaming telemetry reports for Rsyslog servers.

Older versions of rsyslog servers may intermittently miss few report data in some reports. You can upgrade to a newer version to avoid this issue.

Active Directory

Active Directory login failed. How to resolve this?

344 Frequently asked questions


To diagnose the problem, on the Active Directory Configuration and Management page, click Test Settings. Review the test results and fix the problem. Change the configuration and run the test until the test user passes the authorization step.

In general, check the following:

- While logging in, make sure that you use the correct user domain name and not the NetBIOS name. If you have a local iDRAC user account, log into iDRAC using the local credentials. After logging in, make sure that:
 - The Active Directory Enabled option is selected on the Active Directory Configuration and Management page.
 - The DNS setting is correct on the **iDRAC Networking configuration** page.
 - The correct Active Directory root CA certificate is uploaded to iDRAC if certificate validation was enabled.
 - The iDRAC name and iDRAC Domain name matches the Active Directory environment configuration if you are using extended schema.
 - The Group Name and Group Domain Name matches the Active Directory configuration if you are using standard schema.
 - If the user and the iDRAC object is in different domain, then do not select the User Domain from Login option. Instead select **Specify a Domain** option and enter the domain name where the iDRAC object resides.
- Check the domain controller SSL certificates to make sure that the iDRAC time is within the valid period of the certificate.

Active Directory login fails even if certificate validation is enabled. The test results display the following error message. Why does this occur and how to resolve this?

ERROR: Can't contact LDAP server, error:14090086:SSL routines:SSL3 GET SERVER CERTIFICATE:certificate verify failed: Please check the correct Certificate Authority (CA) certificate has been uploaded to iDRAC. Please also check if the iDRAC date is within the valid period of the certificates and if the Domain Controller Address configured in iDRAC matches the subject of the Directory Server Certificate.

If certificate validation is enabled, when iDRAC establishes the SSL connection with the directory server, iDRAC uses the uploaded CA certificate to verify the directory server certificate. The most common reasons for failing certification validation are:

- iDRAC date is not within the validity period of the server certificate or CA certificate. Check the iDRAC time and the validity • period of your certificate.
- The domain controller addresses configured in iDRAC does not match the Subject or Subject Alternative Name of the directory server certificate. If you are using an IP address, read the next question. If you are using FQDN, make sure you are using the FQDN of the domain controller and not the domain. For example, servername.example.com instead of example.com.

Certificate validation fails even if IP address is used as the domain controller address. How to resolve this?

Check the Subject or Subject Alternative Name field of your domain controller certificate. Normally, Active Directory uses the host name and not the IP address of the domain controller in the Subject or Subject Alternative Name field of the domain controller certificate. To resolve this, do any of the following:

- Configure the host name (FQDN) of the domain controller as the domain controller address(es) on iDRAC to match the Subject or Subject Alternative Name of the server certificate.
- Reissue the server certificate to use an IP address in the Subject or Subject Alternative Name field, so that it matches the IP address configured in iDRAC.
- Disable certificate validation if you choose to trust this domain controller without certificate validation during the SSL handshake.

How to configure the domain controller address(es) when using extended schema in a multiple domain environment?

This must be the host name (FQDN) or the IP address of the domain controller(s) that serves the domain in which the iDRAC object resides.

When to configure Global Catalog Address(es)?

If you are using standard schema and the users and role groups are from different domains, Global Catalog Address(es) are required. In this case, you can use only Universal Group.

If you are using standard schema and all the users and role groups are in the same domain, Global Catalog Address(es) are not required.

If you are using extended schema, the Global Catalog Address is not used.

How does standard schema query work?

iDRAC connects to the configured domain controller address(es) first. If the user and role groups are in that domain, the privileges are saved.



If Global Controller Address(es) is configured, iDRAC continues to query the Global Catalog. If additional privileges are retrieved from the Global Catalog, these privileges are accumulated.

Does iDRAC always use LDAP over SSL?

Yes. All the transportation is over secure port 636 and/or 3269. During test setting, iDRAC does a LDAP CONNECT only to isolate the problem, but it does not do an LDAP BIND on an insecure connection.

Why does iDRAC enable certificate validation by default?

iDRAC enforces strong security to ensure the identity of the domain controller that iDRAC connects to. Without certificate validation, a hacker can spoof a domain controller and hijack the SSL connection. If you choose to trust all the domain controllers in your security boundary without certificate validation, you can disable it through the Web interface or RACADM.

Does iDRAC support the NetBIOS name?

Not in this release.

Why does it take up to four minutes to log in to iDRAC using Active Directory Single Sign-On or Smart Card Login?

The Active Directory Single Sign–On or Smart Card log in normally takes less than 10 seconds, but it may take up to four minutes to log in if you have specified the preferred DNS server and the alternate DNS server, and the preferred DNS server has failed. DNS time-outs are expected when a DNS server is down. iDRAC logs you in using the alternate DNS server.

The Active Directory is configured for a domain present in Windows Server 2008 Active Directory. A child or sub domain is present for the domain, the user and group is present in the same child domain, and the user is a member of that group. When trying to log in to iDRAC using the user present in the child domain, Active Directory Single Sign-On login fails.

This may be because of the an incorrect group type. There are two kinds of Group types in the Active Directory server:

- Security Security groups allow you to manage user and computer access to shared resources and to filter group policy settings.
- Distribution Distribution groups are intended to be used only as email distribution lists.

Always make sure that the group type is Security. You cannot use distribution groups to assign permission on any object, however use them to filter group policy settings.

Single Sign-On

SSO login fails on Windows Server 2008 R2 x64. What are the settings required to resolve this?

1. Run the technet.microsoft.com/en-us/library/dd560670(WS.10).aspx for the domain controller and domain policy.

2. Configure the computers to use the DES-CBC-MD5 cipher suite.

These settings may affect compatibility with client computers or services and applications in your environment. The Configure encryption types allowed for Kerberos policy setting is located at **Computer Configuration** > **Security Settings** > **Local Policies** > **Security Options**.

- 3. Make sure that the domain clients have the updated GPO.
- 4. At the command line, type gpupdate /force and delete the old key tab with klist purge command.
- 5. After the GPO is updated, create the new keytab.
- 6. Upload the keytab to iDRAC.

You can now log in to iDRAC using SSO.

Why does SSO login fail with Active Directory users on Windows 7 and Windows Server 2008 R2?

You must enable the encryption types for Windows 7 and Windows Server 2008 R2. To enable the encryption types:

- 1. Log in as administrator or as a user with administrative privilege.
- 2. Go to Start and run gpedit.msc. The Local Group Policy Editor window is displayed.
- 3. Go to Local Computer Settings > Windows Settings > Security Settings > Local Policies > Security Options.
- 4. Right-click Network Security: Configure encryption types allowed for kerberos and select Properties.
- 5. Enable all the options.

6. Click OK. You can now log in to iDRAC using SSO.

Perform the following additional settings for Extended Schema:

- 1. In the Local Group Policy Editor window, navigate to Local Computer Settings > Windows Settings > Security Settings > Local Policies > Security Options .
- 2. Right-click Network Security: Restrict NTLM: Outgoing NTLM traffic to remote server and select Properties.



- 3. Select Allow all, click OK, and close the Local Group Policy Editor window.
- 4. Go to Start and run cmd. The command prompt window is displayed.
- 5. Run the command gpupdate /force. The group policies are updated. Close the command prompt window.
- 6. Go to Start and run regedit. The Registry Editor window is displayed.
- 7. Navigate to $HKEY_LOCAL_MACHINE > System > CurrentControlSet > Control > LSA$.
- 8. In the right-pane, right-click and select New > DWORD (32-bit) Value.
- 9. Name the new key as **SuppressExtendedProtection**.
- 10. Right-click SuppressExtendedProtection and click Modify.
- 11. In the Value data field, type 1 and click OK.
- 12. Close the Registry Editor window. You can now log in to iDRAC using SSO.

If you have enabled SSO for iDRAC and you are using Internet Explorer to log in to iDRAC, SSO fails and you are prompted to enter your user name and password. How to resolve this?

Make sure that the iDRAC IP address is listed in the **Tools** > **Internet Options** > **Security** > **Trusted sites**. If it is not listed, SSO fails and you are prompted to enter your user name and password. Click **Cancel** and proceed.

Smart card login

It takes up to four minutes to log into iDRAC using Active Directory Smart Card login.

The normal Active Directory Smart Card login normally takes less than 10 seconds, however it may take up to four minutes if you have specified the preferred DNS server and the alternate DNS server in the **Network** page, and the preferred DNS server has failed. DNS time-outs are expected when a DNS server is down. iDRAC logs you in using the alternate DNS server.

ActiveX plug-in unable to detect the Smart Card reader.

Make sure that the smart card is supported on the Microsoft Windows operating system. Windows supports a limited number of smart card Cryptographic Service Providers (CSPs).

In general, check if the smart card CSPs are present on a particular client, insert the smart card in the reader at the Windows logon (Ctrl-Alt-Del) screen and check if Windows detects the smart card and displays the PIN dialog-box.

Incorrect Smart Card PIN.

Check if the smart card is locked due to too many attempts with an incorrect PIN. In such cases, contact the smart card issuer in the organization to get a new smart card.

Virtual console

What is the required Java version to launch Virtual Console?

You need Java 8 or later to use this feature and to launch iDRAC Virtual Console over an IPv6 network.

Virtual Console session is active even if you have logged out of iDRAC web interface. Is this the expected behavior?

Yes. Close the Virtual Console Viewer window to log out of the corresponding session.

Can a new remote console video session be started when the local video on the server is turned off?

Yes.

Why does it take 15 seconds to turn off the local video on the server after requesting to turn off the local video?

It gives a local user an opportunity to take any action before the video is switched off.

Is there a time delay when turning on the local video?

No, after a local video turn ON request is received by iDRAC, the video is turned on instantly.

Can the local user also turn off or turn on the video?

When the local console is disabled, the local user cannot turn off or turn on the video.

Does switching off the local video also switch off the local keyboard and mouse?

No.

Does turning off the local console turn off the video on the remote console session?

No, turning the local video on or off is independent of the remote console session.



What privileges are required for an iDRAC user to turn on or turn off the local server video?

Any user with iDRAC configuration privileges can turn on or turn off the local console.

How to get the current status of the local server video?

The status is displayed on the Virtual Console page.

To display the status of the object iDRAC.VirtualConsole.AttachState, use the following command:

racadm get idrac.virtualconsole.attachstate

Or, use the following command from a SSH or a remote session:

racadm -r (iDrac IP) -u (username) -p (password) get iDRAC.VirtualConsole.AttachState

The status is also seen on the Virtual Console OSCAR display. When the local console is enabled, a green status is displayed next to the server name. When disabled, a yellow dot indicates that iDRAC has locked the local console.

Why is the bottom of the system screen not seen from the Virtual Console window?

Make sure that the management station's monitor resolution is set to 1280 x 1024.

Why is the Virtual Console Viewer window garbled on Linux operating system?

The console viewer on Linux requires a UTF-8 character set. Check your locale and reset the character set if required.

Why does the mouse not synchronize under the Linux text console in Lifecycle Controller?

Virtual Console requires the USB mouse driver, but the USB mouse driver is available only under the X-Window operating system. In the Virtual Console viewer, do any of the following:

- Go to Tools > Session Options > Mouse tab. Under Mouse Acceleration, select Linux.
- Under the Tools menu, select Single Cursor option.

How to synchronize the mouse pointers on the Virtual Console Viewer window?

Before starting a Virtual Console session, make sure that the correct mouse is selected for your operating system.

Make sure that the **Single Cursor** option under **Tools** in the iDRAC Virtual Console menu is selected on iDRAC Virtual Console client. The default is two cursor mode.

Can a keyboard or mouse be used while installing a Microsoft operating system remotely through the Virtual Console?

No. When you remotely install a supported Microsoft operating system on a system with Virtual Console enabled in the BIOS, an EMS Connection Message is sent that requires that you select **OK** remotely. You must either select **OK** on the local system or restart the remotely managed server, reinstall, and then turn off the Virtual Console in BIOS.

This message is generated by Microsoft to alert the user that Virtual Console is enabled. To make sure that this message does not appear, always turn off Virtual Console in the iDRAC Settings utility before remotely installing an operating system.

Why does the Num Lock indicator on the management station not reflect the status of the Num Lock on the remote server?

When accessed through the iDRAC, the Num Lock indicator on the management station does not necessarily coincide with the state of the Num Lock on the remote server. The state of the Num Lock depends the setting on the remote server when the remote session is connected, regardless of the state of the Num Lock on the management station.

Why do multiple Session Viewer windows appear when a Virtual Console session is established from the local host?

You are configuring a Virtual Console session from the local system. This is not supported.

If a Virtual Console session is in-progress and a local user accesses the managed server, does the first user receive a warning message?

No. If a local user accesses the system, both have control of the system.

How much bandwidth is required to run a Virtual Console session?

It is recommended to have a 5 MBPS connection for good performance. A 1 MBPS connection is required for minimal performance.

What is the minimum system requirements for the management station to run Virtual Console?

The management station requires an Intel Pentium III 500 MHz processor with at least 256 MB of RAM.

Why does Virtual Console Viewer window sometimes displays No Signal message?



You may see this message because the iDRAC Virtual Console plug-in is not receiving the remote server desktop video. Generally, this behavior may occur when the remote server is turned off. Occasionally, the message may be displayed due to a remote server desktop video reception malfunction.

Why does Virtual Console Viewer window sometimes display an Out of Range message?

You may see this message because a parameter necessary to capture video is beyond the range for which the iDRAC can capture the video. Parameters such as display resolution or refresh rate too high causes an out of range condition. Normally, physical limitations such as video memory size or bandwidth sets the maximum range of parameters.

When starting a Virtual Console session from iDRAC web interface, why is an ActiveX security popup displayed?

iDRAC may not be in the trusted site list. To prevent the security popup from appearing every time you begin a Virtual Console session, add iDRAC to the trusted site list in the client browser:

- 1. Click Tools > Internet Options > Security > Trusted sites.
- 2. Click Sites and enter the IP address or the DNS name of iDRAC
- 3. Click Add.
- 4. Click Custom Level.
- 5. In the Security Settings window, select Prompt under Download unsigned ActiveX Controls.

Why is the Virtual Console Viewer window blank?

If you have Virtual Media privilege, but not Virtual Console privilege, you can start the viewer to access the virtual media feature, but the managed server's console is not displayed.

Why doesn't the mouse synchronize in DOS when using Virtual Console?

The Dell BIOS is emulating the mouse driver as a PS/2 mouse. By design, the PS/2 mouse uses relative position for the mouse pointer, which causes the lag in syncing. iDRAC has a USB mouse driver that allows absolute position and closer tracking of the mouse pointer. Even if iDRAC passes the USB absolute mouse position to the Dell BIOS, the BIOS emulation converts it back to relative position and the behavior remains. To fix this problem, set the mouse mode to USC/Diags in the Configuration screen.

After launching the Virtual Console, the mouse cursor is active on the Virtual Console, but not on the local system. Why does this occur and how to resolve this?

This occurs if the **Mouse Mode** is set to **USC/Diags**. Press **Alt + M** hot key to use the mouse on the local system. Press **Alt + M** again to use the mouse on the Virtual Console.

Why does GUI session time out after launching a virtual console from the iDRAC interface that is launched from CMC?

When launching the Virtual Console to iDRAC from the CMC web interface a popup is opened to launch the Virtual Console. The popup closes shortly after the Virtual Console opens.

When launching both the GUI and Virtual Console to the same iDRAC system on a management station, a session time-out for the iDRAC GUI occurs if the GUI is launched before the popup closes. If the iDRAC GUI is launched from the CMC web interface after the popup with the Virtual Console closed, this issue does not appear.

(i) NOTE: Not applicable for MX platforms.

Why does Linux SysRq key not work with Internet Explorer?

The Linux SysRq key behavior is different when using Virtual Console from Internet Explorer. To send the SysRq key, press the **Print Screen** key and release while holding the **Ctrl** and **Alt** keys. To send the SysRq key to a remote Linux server though iDRAC, while using Internet Explorer:

1. Activate the magic key function on the remote Linux server. You can use the following command to activate it on the Linux terminal:

echo 1 > /proc/sys/kernel/sysrq

- 2. Activate the keyboard pass-through mode of Active X Viewer.
- 3. Press Ctrl+Alt+Print Screen.
- 4. Release only Print Screen.
- 5. Press Print Screen+Ctrl+Alt.

(i) **NOTE:** The SysRq feature is currently not supported with Internet Explorer and Java.

Why is the "Link Interrupted" message displayed at the bottom of the Virtual Console?

When using the shared network port during a server reboot, iDRAC is disconnected while BIOS is resetting the network card. This duration is longer on 10 Gb cards, and is also exceptionally long if the connected network switch has Spanning Tree



Protocol (STP) enabled. In this case, it is recommended to enable "portfast" for the switch port connected to the server. In most cases, the Virtual Console restores itself.

Launching Virtual Console with Java plug-in fails after the iDRAC firmware was updated.

Delete the Java cache and then launch the virtual console.

To enable console redirection using the web server port (443)

racadm>>set iDRAC.VirtualConsole.WebRedirect Enabled

To close the external virtual console port (5900), set the following iDRAC property.

To close the external virtual console port (5900), both iDRAC.VirtualConsole.WebRedirect and iDRAC.VirtualConsole.CloseUnusedPort must be enabled.

racadm>>set iDRAC.VirtualConsole.CloseUnusedPort Enabled

(i) NOTE:

- If the Virtual Media port is disabled, the stand-alone Virtual Media will not be accessible and you may use the Virtual Media through Virtual Console.
- While CloseUnusedPort is enabled, Java and ActiveX based Virtual console and Virtual media will not function as they
 require dedicated external port. Virtual console and Virtual media using HTML5 plug-in will function on iDRAC web
 server port (443).

Virtual media

Why does the Virtual Media client connection sometimes drop?

When a network time-out occurs, iDRAC firmware drops the connection, disconnecting the link between the server and the virtual drive.

If you change the CD in the client system, the new CD may have an autostart feature. In this case, the firmware can time out and the connection is lost if the client system takes too long to read the CD. If a connection is lost, reconnect from the GUI and continue the previous operation.

If the Virtual Media configuration settings are changed in the iDRAC web interface or through local RACADM commands, any connected media is disconnected when the configuration change is applied.

To reconnect to the Virtual Drive, use the Virtual Media Client View window.

Why does a Windows operating system installation through Virtual Media take an extended amount of time?

If you are installing the Windows operating system using the *Dell Systems Management Tools and Documentation DVD* and the network connection is slow, the installation procedure may require an extended amount of time to access iDRAC web interface due to network latency. The installation window does not indicate the installation progress.

How to configure the virtual device as a bootable device?

On the managed system, access BIOS Setup and go to the boot menu. Locate the virtual CD, virtual floppy, or vFlash and change the device boot order as required. Also, press the "spacebar" key in the boot sequence in the CMOS setup to make the virtual device bootable. For example, to boot from a CD drive, configure the CD drive as the first device in the boot order.

What are the types of media that can be set as a bootable device?

iDRAC allows you to boot from the following bootable media:

- CDROM/DVD Data media
- ISO 9660 image
- 1.44 Floppy disk or floppy image
- A USB key that is recognized by the operating system as a removable disk
- A USB key image

How to make the USB key a bootable device?

You can also boot with a Windows 98 startup disk and copy system files from the startup disk to the USB key. For example, from the DOS prompt, type the following command:

sys a: x: /s



where, x: is the USB key that is required to be set as a bootable device.

The Virtual Media is attached and connected to the remote floppy. But, cannot locate the Virtual Floppy/Virtual CD device on a system running Red Hat Enterprise Linux or the SUSE Linux operating system. How to resolve this?

Some Linux versions do not auto-mount the virtual floppy drive and the virtual CD drive in the same method. To mount the virtual floppy drive, locate the device node that Linux assigns to the virtual floppy drive. To mount the virtual floppy drive:

1. Open a Linux command prompt and run the following command:

grep "Virtual Floppy" /var/log/messages

- 2. Locate the last entry to that message and note the time.
- 3. At the Linux prompt, run the following command:

grep "hh:mm:ss" /var/log/messages

where, hh:mm:ss is the time stamp of the message returned by grep in step 1.

- 4. In step 3, read the result of the grep command and locate the device name that is given to the Virtual Floppy.
- 5. Make sure that you are attached and connected to the virtual floppy drive.
- 6. At the Linux prompt, run the following command:

mount /dev/sdx /mnt/floppy

where, /dev/sdx is the device name found in step 4 and /mnt/floppy is the mount point.

- To mount the virtual CD drive, locate the device node that Linux assigns to the virtual CD drive. To mount the virtual CD drive:
- 1. Open a Linux command prompt and run the following command:

grep "Virtual CD" /var/log/messages

- 2. Locate the last entry to that message and note the time.
- 3. At the Linux prompt, run the following command:

grep "hh:mm:ss" /var/log/messages

where, hh:mm:ss is the timestamp of the message returned by grep in step 1.

- 4. In step 3, read the result of the grep command and locate the device name that is given to the Dell Virtual CD.
- **5.** Make sure that the Virtual CD Drive is attached and connected.
- 6. At the Linux prompt, run the following command:

mount /dev/sdx /mnt/CD

where: /dev/sdx is the device name found in step 4 and /mnt/floppy is the mount point.

Why are the virtual drives attached to the server removed after performing a remote firmware update using the iDRAC web interface?

Firmware updates cause the iDRAC to reset, drop the remote connection, and unmount the virtual drives. The drives reappear when iDRAC reset is complete.

Why are all the USB devices detached after connecting a USB device?

Virtual media devices and vFlash devices are connected as a composite USB device to the Host USB BUS, and they share a common USB port. Whenever any virtual media or vFlash USB device is connected to or disconnected from the host USB bus, all the Virtual Media and vFlash devices are disconnected momentarily from the host USB bus, and then they are reconnected. If the host operating system uses a virtual media device, do not attach or detach one or more virtual media or vFlash devices. It is recommended that you connect all the required USB devices first before using them.

What does the USB Reset do?

It resets the remote and local USB devices connected to the server.

How to maximize Virtual Media performance?

To maximize Virtual Media performance, launch the Virtual Media with the Virtual Console disabled or do one of the following:

- Change the performance slider to Maximum Speed.
- Disable encryption for both Virtual Media and Virtual Console.



NOTE: In this case, the data transfer between managed server and iDRAC for Virtual Media and Virtual Console will not be secured.

If you are using any Windows server operating systems, stop the Windows service named Windows Event Collector. To do
this, go to Start > Administrative Tools > Services. Right-click Windows Event Collector and click Stop.

While viewing the contents of a floppy drive or USB key, a connection failure message is displayed if the same drive is attached through the virtual media?

Simultaneous access to virtual floppy drives is not allowed. Close the application used to view the drive contents before attempting to virtualize the drive.

What file system types are supported on the Virtual Floppy Drive?

The virtual floppy drive supports FAT16 or FAT32 file systems.

Why is an error message displayed when trying to connect a DVD/USB through virtual media even though the virtual media is currently not in use?

The error message is displayed if Remote File Share (RFS) feature is also in use. At a time, you can use RFS or Virtual Media and not both.

Virtual Media inaccessible even though iDRAC shows Virtual Media Connection Status as Connected.

If you try to access Virtual Media using ActiveX or Java plug-in while **Attach Mode** is set to **Detach** in iDRAC, then the connection status may show as **Connected**. Change the **Attach Mode** to either **Auto-attach** or **Attach** to access the Virtual Media.

vFlash SD card

When is the vFlash SD card locked?

The vFlash SD card is locked when an operation is in-progress. For example, during an initialize operation.

SNMP authentication

Why is the message 'Remote Access: SNMP Authentication Failure' displayed?

As part of discovery, IT Assistant attempts to verify the get and set community names of the device. In IT Assistant, you have the get community name = public and the set community name = private. By default, the SNMP agent community name for iDRAC agent is public. When IT Assistant sends out a set request, the iDRAC agent generates the SNMP authentication error because it accepts requests only from community = public.

To prevent SNMP authentication errors from being generated, you must enter community names that are accepted by the agent. Since the iDRAC only allows one community name, you must use the same get and set community name for IT Assistant discovery setup.

Storage devices

OpenManage Storage Management displays more storage devices than iDRAC while information for all the storage devices connected to the system are not displayed. Why?

iDRAC displays information for only the Comprehensive Embedded Management (CEM) supported devices.

For External JBODs/Insights behind the HBA, EEMI Message for the SAS Connector/IOM removal is generated with the EEMI message ID ENC42, however, EEMI message ENC41 for the SAS Connector/IOM restoration is not generated.

To confirm restoration of the IOM in iDRAC web interface:

- 1. Go to Storage > Overview > Enclosures
- 2. Select the enclosure.
- 3. Under Advanced Properties, ensure that the value for Redundant Path is set to Present, then IOM restoration is confirmed.



GPU (Accelerators)

Accelerators section under CPU/Accelerators in iDRAC GUI is grayed out.

Few pages in GUI may not show expected response when respective attribute is disabled in Redfish.

iDRAC Service Module

iSM details are missing / not updated correctly in iDRAC GUI page of some PowerEdge servers

When a user adds SUB NIC under teaming, the configuration is invalid. This causes iSM to not to communicate with iDRAC properly.

Before installing or running the iDRAC Service Module, should the OpenManage Server Administrator be uninstalled?

No you do not have to uninstall Server Administrator. Before you install or run the iDRAC Service Module, make sure that you have stopped the features of Server Administrator that the iDRAC Service Module provides.

How to check whether iDRAC Service Module is installed in the host operating system?

To know if the iDRAC Service Module is installed on the system,

On systems running Windows:

Open the **Control Panel**, verify if iDRAC Service Module is listed in the list of installed programs displayed.

• On systems running Linux:

Run the command rpm -qi dcism. If the iDRAC Service Module is installed, the status displayed is **installed**.

- On systems running ESXi: Run the command esxcli software vib list|grep -i open on the host. iDRAC Service module is displayed.
- () NOTE: To check if the iDRAC Service Module is installed on Red Hat Enterprise Linux 7, use the systemctl status dcismeng.service command instead of the init.d command.

How to check the version number of the iDRAC Service Module installed in the system?

To check the version of the iDRAC Service Module in the system, do any of the following:

- Click Start > Control Panel > Programs and Features. The version of the installed iDRAC Service Module is listed in the Version tab.
- Go to My Computer > Uninstall or change a program.

What is the minimum permission level required to install the iDRAC Service Module?

To install the iDRAC Service Module, you must have administrator level privileges.

On iDRAC Service Module version 2.0 and earlier, while installing the iDRAC Service Module, an error message is displayed stating this is not a supported server. Consult the User Guide for additional information about the supported servers. How to resolve this error?

Before installing the iDRAC Service Module, make sure that the server is a 12th generation PowerEdge server or later. Also, make sure that you have a 64-bit system.

The following message is displayed in the OS log, even when the OS to iDRAC Pass-through over USBNIC is configured properly. Why?

The iDRAC Service Module is unable to communicate with iDRAC using the OS to iDRAC Pass-through channel

iDRAC Service Module uses the OS to iDRAC pass-through over USB NIC feature to establish the communication with iDRAC. Sometimes, the communication is not established though the USB NIC interface is configured with the correct IP endpoints. This may happen when the host operating system routing table has multiple entries for the same destination mask and the USB NIC destination is not listed as the first one in routing order.

Table 64. Example of a routing order

Destination	Gateway	Genmask	Flags	Metric	Ref	Use Iface
default	10.94.148.1	0.0.0.0	UG	1024	0	0 em1
10.94.148.0	0.0.0.0	255.255.255.0	U	0	0	0 em1



Table 64. Example of a routing order (continued)

Destination	Gateway	Genmask	Flags	Metric	Ref	Use Iface
link-local	0.0.0.0	255.255.255.0	U	0	0	0 em1
link-local	0.0.0.0	255.255.255.0	U	0	0	0 enp0s20u12u3

In the example **enp0s20u12u3** is the USB NIC interface. The link-local destination mask is repeated and the USB NIC is not the first one in the order. This results in the connectivity issue between iDRAC Service Module and iDRAC over the OS to iDRAC Pass-through. To troubleshoot the connectivity issue, make sure that the iDRAC USBNIC IPv4 address (by default it is 169.254.1.1) is reachable from the host operating system.

If not:

- Change the iDRAC USBNIC address on a unique destination mask.
- Delete the entries that are not required from the routing table to make sure that USB NIC is chosen by route when the host wants to reach the iDRAC USB NIC IPv4 address.

On iDRAC Service Module version 2.0 and earlier, when uninstalling iDRAC Service Module from a VMware ESXi server, the virtual switch is named as vSwitchiDRACvusb and port group as iDRAC Network on the vSphere client. How to delete them?

While installing iDRAC Service Module VIB on a VMware ESXi server, iDRAC Service Module creates the vSwtich and Portgroup to communicate with iDRAC over the OS to iDRAC Pass-through in USB NIC mode. After the uninstallation, the virtual switch **vSwitchiDRACvusb** and the port group **iDRAC Network** are not deleted. To delete it manually, perform one of the following steps:

- Go to vSphere Client Configuration wizard and delete the entries.
- Go to the Esxcli and type the following commands:
 - To remove port group: esxcfg-vmknic -d -p "iDRAC Network"
 - To remove vSwitch: esxcfg-vswitch -d vSwitchiDRACvusb
 - **NOTE:** You can reinstall iDRAC Service Module on the VMware ESXi server as this is not a functional issue for the server.

Where is the Replicated Lifecycle log available on the operating system?

To view the replicated Lifecycle logs:

Table 65. Lifecycle logs location

Operating System	Location		
Microsoft Windows	 Event viewer > Windows Logs > System. All the iDRAC Service Module Lifecycle logs are replicated under the source name iDRAC Service Module. (i) NOTE: In iSM version 2.1 and later, Lifecycle logs are replicated under the Lifecycle Controller Log source name. In iSM version 2.0 and earlier, the logs are replicated under iDRAC Service Module source name. (i) NOTE: The location of the Lifecycle log can be configured using the iDRAC Service Module installer. You can configure the location while installing iDRAC Service Module or modifying the installer. 		
Red Hat Enterprise Linux, SUSE Linux, CentOS, and Citrix XenServer	/var/log/messages		
VMware ESXi	/var/log/syslog.log		

What are the Linux-dependent packages or executables available for installation while completing the Linux installation?

To see the list of Linux-dependent packages, see the *Linux Dependencies* section in the *iDRAC Service Module User's Guide* available at https://www.dell.com/idracmanuals.

How to increase GPU performance for certain configuration?

BIOS system performance profile set to performance



Under Processor Settings, set NPS to 4 and CCX to auto Minimum 1 DIMM per channel IOmmu=passthrough on Linux OS

RACADM

After performing an iDRAC reset (using the racadm racreset command), if any command is issued, the following message is displayed. What does this indicate?

ERROR: Unable to connect to RAC at specified IP address

The message indicates that you must wait until the iDRAC completes the reset before issuing another command.

When using RACADM commands and subcommands, some errors are not clear.

- You may see one or more of the following errors when using the RACADM commands:
- Local RACADM error messages Problems such as syntax, typographical errors, and incorrect names.
- Remote RACADM error messages Problems such as incorrect IP Address, incorrect user name, or incorrect password.

During a ping test to iDRAC, if the network mode is switched between Dedicated and Shared modes, there is no ping response.

Clear the ARP table on your system.

Remote RACADM fails to connect to iDRAC from SUSE Linux Enterprise Server (SLES) 11 SP1.

Make sure that the official openssl and libopenssl versions are installed. Run the following command to install the RPM packages:

rpm -ivh --force < filename >

where, filename is the openssl or libopenssl rpm package file.

For example:

rpm -ivh --force openssl-0.9.8h-30.22.21.1.x86_64.rpm rpm -ivh --force libopenssl0_9_8-0.9.8h-30.22.21.1.x86_64.rpm

Why are the remote RACADM and web-based services unavailable after a property change?

It may take a while for the remote RACADM services and the Web-based interface to become available after the iDRAC web server resets.

The iDRAC Web server is reset when:

- The network configuration or network security properties are changed using the iDRAC web user interface.
- The iDRAC.Webserver.HttpsPort property is changed, including when a racadm set -f <config file> changes it.
- The racresetcfg command is used.
- iDRAC is reset.
- A new SSL server certificate is uploaded.

Why is an error message displayed if you try to delete a partition after creating it using local RACADM?

This occurs because the create partition operation is in-progress. However, the partition is deleted after sometime and a message that the partition is deleted is displayed. If not, wait until the create partition operation is completed and then delete the partition.

Permanently setting the default password to calvin

If your system shipped with a unique default iDRAC password but you want to set *calvin* as the default password, you must use the jumpers available on the system board.

CAUTION: Changing the jumper settings permanently changes the default password to *calvin.* You cannot revert to the unique password even if you reset iDRAC to factory settings.

Frequently asked questions 355

For information about the jumper location and the procedure, see the documentation for your server at https://www.dell.com/support.

Fls.<u>1398</u> Mov. 33



Upgrade fails when upgrading to the latest version.

(i) NOTE: 3.30.30.30 is the minimum iDRAC version required to upgrade to 4.00.00.00 / 4.10.10.10 of later build .

After an iDRAC reset, iDRAC GUI may not display all the values.

NOTE: If you reset the iDRAC for some reason, ensure that you wait for at least two minutes after resetting iDRAC to access or modify any settings in iDRAC.

When an OS is installed, hostname may or may not appear/change automatically.

There are two scenarios:

- Scenario 1: iDRAC is not showing the latest hostname once you install an OS. You need to install OMSA or iSM along with the iDRAC to get the hostname reflected.
- Scenario 2: iDRAC had a hostname for a specific OS and another different OS has been installed and still the hostname is appearing as the old hostname without overwriting the hostname. The reason behind, hostname is an information which is coming from the OS, iDRAC only saves the information. If there is a new OS has been installed, iDRAC does not reset the value of the hostname. However, newer versions of the OSs are capable to update the hostname in iDRAC during the 1st OS startup.

How to find an iDRAC IP address for a blade server?

(i) NOTE: The Chassis Management Controller (CMC) option is applicable only for Blade servers.

• Using CMC web interface:

Go to **Chassis** > **Servers** > **Setup** > **Deploy**. In the table that is displayed, view the IP address for the server.

• Using the Virtual Console: Reboot the server to view the iDRAC IP address during POST. Select the "Dell CMC" console in the OSCAR interface to log in to CMC through a local serial connection. CMC RACADM commands can be sent from this connection.

For more information on CMC RACADM commands, see the *Chassis Management Controller RACADM CLI Guide* available at https://www.dell.com/cmcmanuals.

For more information on iDRAC RACADM commands, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/ idracmanuals.

• Using local RACADM

Use the command: racadm getsysinfo For example:

```
$ racadm getniccfg -m server-1
DHCP Enabled = 1
IP Address = 192.168.0.1
Subnet Mask = 255.255.255.0
Gateway = 192.168.0.1
```

• Using LCD:

On the main menu, highlight the server, press the check button, select the required server, and press the check button.



How to find an iDRAC IP address for a blade server ?

(i) NOTE: The OME-Modular web interface option is applicable only for MX platforms.

• Using OME-Modular web interface:

Go to **Devices** > **Compute**. Select the computer sled and iDRAC IP is displayed as **Management IP**.

- Using OMM Application: see the Dell EMC OpenManage Mobile User's Guide available at https://www.dell.com/ openmanagemanuals
- Using Serial Connection
- Using LCD: On the main menu, highlight the server, press the check button, select the required server, and press the check button.

How to find the CMC IP address related to the blade server?

(i) NOTE: Not applicable for MX platforms.

• From iDRAC web interface:

Go to **iDRAC Settings** > **CMC**. The **CMC Summary** page displays the CMC IP address.

• From the Virtual Console:

Select the "Dell CMC" console in the OSCAR interface to log in to CMC through a local serial connection. CMC RACADM commands can be issued from this connection.

```
$ racadm getniccfg -m chassis
NIC Enabled
                    = 1
                   = 1
DHCP Enabled
                   = 192.168.0.120
Static IP Address
Static Subnet Mask = 255.255.255.0
                    = 192.168.0.1
Static Gateway
Current IP Address = 10.35.155.151
Current Subnet Mask = 255.255.255.0
Current Gateway
                   = 10.35.155.1
Speed
                    = Autonegotiate
Duplex
                    = Autonegotiate
```

(i) NOTE: You can also perform this using remote RACADM.

For more information on CMC RACADM commands, see the *Chassis Management Controller RACADM CLI Guide* available at https://www.dell.com/cmcmanuals.

For more information on iDRAC RACADM commands, see the *iDRAC RACADM CLI Guide* available at https://www.dell.com/ idracmanuals.

How to find the OME Modular IP address?

(i) NOTE: Applicable only for MX platforms.

• From iDRAC web interface:

Go to **iDRAC Settings** > Management Module. The Management Module page displays the OME Modular IP address.

How to find iDRAC IP address for rack and tower server?

- From Local RACADM:
 - Use the command racadm getsysinfo.
- From LCD:



On the physical server, use the LCD panel navigation buttons to view the iDRAC IP address. Go to **Setup View** > **View** > **iDRAC IP** > **IPv4** or **IPv6** > **IP**.

From OpenManage Server Administrator:

In the Server Administrator web interface, go to Modular Enclosure > System/Server Module > Main System Chassis/ Main System > Remote Access.

iDRAC network connection is not working.

For blade servers:

- Ensure that the LAN cable is connected to CMC. (not for MX platforms)
- Ensure that NIC settings, IPv4 or IPv6 settings, and either Static or DHCP is enabled for your network.

For rack and tower servers:

- In shared mode, ensure that the LAN cable is connected to the NIC port where the wrench symbol is present.
- In Dedicated mode, ensure that the LAN cable is connected to the iDRAC LAN port.
- Ensure that NIC settings, IPv4 and IPv6 settings and either Static or DHCP is enabled for your network.

iDRAC not accessible in shared LOM

iDRAC may be inaccessible if there are fatal errors in host OS such as BSOD error in Windows. To access iDRAC, reboot the host to recover the connection.

Shared LOM not functional after enabling Link Aggregation Control Protocol (LACP).

The host OS driver for the network adapter must be loaded before LACP is enabled. However, if a passive LACP configuration is in use, the shared LOM may be functional before the host OS driver is loaded. See the switch documentation for LACP configuration.

(i) NOTE: Shared LOM IP of iDRAC is not accessible in pre-boot state when the switch is configured with LACP.

Inserted the blade server into the chassis and pressed the power switch, but it did not power on.

- iDRAC requires up to two minutes to initialize before the server can power on.
- Check CMC, andOME Modular (only for MX platforms) power budget. The chassis power budget may have exceeded.

How to retrieve an iDRAC administrative user name and password?

You must restore iDRAC to its default settings. For more information, see Resetting iDRAC to factory default settings.

How to change the name of the slot for the system in a chassis?

(i) NOTE: Not applicable for MX platforms.

- 1. Log in to CMC web interface and go to Chassis > Servers > Setup.
- 2. Enter the new name for the slot in the row for your server and click Apply.

iDRAC on blade server is not responding during boot.

Remove and reinsert the server.

358 Frequently asked questions



Check CMC (not for MX platforms), and OME Modular (Applicable for MX platforms) web interface to see if iDRAC is displayed as an upgradable component. If it does, follow the instructions in Updating firmware using CMC web interface update the firmware.

(i) NOTE: Update feature not applicable for MX platforms.

If the problem persists, contact technical support.

When attempting to boot the managed server, the power indicator is green, but there is no POST or no video.

This happens due to any of the following conditions:

- Memory is not installed or is inaccessible.
- CPU is not installed or is inaccessible
- Video riser card is missing or not connected properly.

Also, see error messages in iDRAC log using iDRAC web interface or from the server LCD.

Unable to login to iDRAC web interface using Firefox browser on Linux or Ubuntu. Unable to enter the password.

To resolve this issue, reinstall or upgrade the Firefox browser.

Unable to access iDRAC through USB NIC in SLES and Ubuntu

(i) NOTE: In SLES, set the iDRAC interface to DHCP.

In Ubuntu, use the Netplan utility to configure iDRAC interface into DHCP mode. To configure the DHCP:

- 1. Use /etc/netplan/01-netcfg.yaml.
- **2.** Specify Yes for iDRAC DHCP.
- **3.** Apply the configuration.





Figure 5. Configuring iDRAC interface to DHCP mode in Ubuntu

Model, Manufacturer and other properties are not listing for Embedded Network Adapters in Redfish

FRU details for embedded devices will not be displayed. There will not be any FRU object for devices which are embedded on Motherboard. Hence dependent property will not be there.



Use case scenarios

This section helps you in navigating to specific sections in the guide to perform typical use case scenarios.

Topics:

- Troubleshooting an inaccessible managed system
- Obtaining system information and assess system health
- Setting up alerts and configuring email alerts
- Viewing and exporting System Event Log and Lifecycle Log
- Interfaces to update iDRAC firmware
- Performing graceful shutdown
- Creating new administrator user account
- Launching servers remote console and mounting a USB drive
- Installing bare metal OS using attached virtual media and remote file share
- Managing rack density
- Installing new electronic license
- Applying IO Identity configuration settings for multiple network cards in single host system reboot

Troubleshooting an inaccessible managed system

After receiving alerts from OpenManage Essentials, Dell Management Console, or a local trap collector, five servers in a data center are not accessible with issues such as hanging operating system or server. Need to identify the cause to troubleshoot and bring up the server using iDRAC.

Before troubleshooting the inaccessible system, make sure that the following prerequisites are met:

- Enable last crash screen
- Alerts are enabled on iDRAC

To identify the cause, check the following in the iDRAC web interface and re-establish the connection to the system:

NOTE: If you cannot access the iDRAC web interface, go to the server, access the LCD panel, write down the IP address or the host name, and then perform the following operations using iDRAC web interface from your management station:

- Server's LED status Blinking amber or Solid amber.
- Front Panel LCD status or error message Amber LCD or error message.
- Operating system image is seen in the Virtual Console. If you can see the image, reset the system (warm boot) and log in again. If you are able to log in, the issue is fixed.
- Last crash screen.
- Boot capture video.
- Crash capture video.
- Server Health status Red *x* icons for the system components with issues.
- Storage array status Possible array offline or failed
- Lifecycle log for critical events related to system hardware and firmware and the log entries that were logged at the time of system crash.
- Generate Tech Support report and view the collected data.
- Use the monitoring features provided by iDRAC Service Module

Obtaining system information and assess system health

To obtain system information and assess system health:



- In iDRAC Web interface, go to **Overview** > **Summary** to view the system information and access various links on this page to asses system health. For example, you can check the health of the chassis fan.
- You can also configure the chassis locator LED and based on the color, assess the system health.
- If iDRAC Service Module is installed, the operating system host information is displayed.

Setting up alerts and configuring email alerts

To set up alerts and configure email alerts:

- 1. Enable alerts.
- 2. Configure the email alert and check the ports.
- 3. Perform a reboot, power off, or power cycle the managed system.
- 4. Send test alert.

Viewing and exporting System Event Log and Lifecycle Log

To view and export lifecycle log and system event log (SEL):

1. In iDRAC Web interface, go to Maintenance > System Event Logs to view SEL and Lifecycle Log to view lifecycle log.

(i) NOTE: The SEL is also recorded in the lifecycle log. Using the filtering options to view the SEL.

- 2. Export the SEL or lifecycle log in the XML format to an external location (management station, USB, network share, and so on). Alternatively, you can enable remote system logging, so that all the logs written to the lifecycle log are also simultaneously written to the configured remote server(s).
- **3.** If you are using the iDRAC Service Module, export the Lifecycle log to OS log.

Interfaces to update iDRAC firmware

Use the following interfaces to update the iDRAC firmware:

- iDRAC Web interface
- Redfish API
- RACADM CLI (iDRAC_) and CMC (not applicable for MX platforms))
- Dell Update Package (DUP)
- CMC (not applicable for MX platforms)OME Modular (applicable only for MX platforms) Web interface
- Lifecycle Controller-Remote Services
- Lifecycle Controller
- Dell Remote Access Configuration Tool (DRACT)

Performing graceful shutdown

To perform graceful shutdown, in iDRAC Web interface, go to one of the following locations:

At Dashboard select Graceful Shutdown and click Apply.

For more information, see the *iDRAC Online Help*.

Creating new administrator user account

You can modify the default local administrator user account or create a new administrator user account. To modify the local administrator user account, see Modifying local administrator account settings.

To create a new administrator account, see the following sections:

• Configuring local users



- Configuring active directory users
- Configuring generic LDAP users

Launching servers remote console and mounting a USB drive

To launch the remote console and mount a USB drive:

- 1. Connect a USB flash drive (with the required image) to the management station.
- 2. Use the following method to launch virtual console through the iDRAC Web Interface:
- Go to Dashboard > Virtual Console and click Launch Virtual Console.

The Virtual Console Viewer is displayed.

- 3. From the File menu, click Virtual Media > Launch Virtual Media.
- Click Add Image and select the image that is located on the USB flash drive. The image is added to the list of available drives.
- 5. Select the drive to map it. The image on the USB flash drive is mapped to the managed system.

Installing bare metal OS using attached virtual media and remote file share

See the Deploying operating system using remote file share section.

Managing rack density

Before you install additional servers in a rack, you must determine the remaining capacity in the rack.

To assess the capacity of a rack to add additional servers:

- 1. View the current power consumption data and historical power consumption data for the servers.
- 2. Based on the data, power infrastructure and cooling system limitations, enable the power cap policy and set the power cap values.
 - **NOTE:** It is recommended that you set a cap close to the peak, and then use that capped level to determine how much capacity is remaining in the rack for adding more servers.

Installing new electronic license

See License operations for more information.

Applying IO Identity configuration settings for multiple network cards in single host system reboot

If you have multiple network cards in a server that is part of a Storage Area Network (SAN) environment and you want to apply different virtual addresses, initiator and target configuration settings to those cards, use the I/O Identity Optimization feature to reduce the time in configuring the settings. To do this:

- 1. Make sure that BIOS, iDRAC, and the network cards are updated to the latest firmware version.
- 2. Enable IO Identity Optimization.
- 3. Export the Server Configuration Profile (SCP) file from iDRAC.
- 4. Edit the I/O Identity optimization settings in the SCP file.

5. Import the SCP file to iDRAC.



Inserido ao protocolo **22.951.206-4** por: **Pedro Henrique Golin Linhares** em: 31/03/2025 12:49. A autenticidade deste documento pode ser validada no endereço: https://www.eprotocolo.pr.gov.br/spiweb/validarDocumento com o código: fc1c44dfe0686d5547ffd691199be699.



Dell PowerEdge RAID Controller 11 User's Guide

PERC H755, H750, H355, and H350 Controller Series

Regulatory Model: UCPA-1101, UCPF-1100, UCPF-1110, UCPN-1100, UCSA-1111, UCSF-1100, and UCSM-1100 September 2023 Rev. A05



Inserido ao protocolo **22.951.206-4** por: **Pedro Henrique Golin Linhares** em: 31/03/2025 12:49. A autenticidade deste documento pode ser validada no endereço: https://www.eprotocolo.pr.gov.br/spiweb/validarDocumento com o código: fc1c44dfe0686d5547ffd691199be699.



Notes, cautions, and warnings

(i) NOTE: A NOTE indicates important information that helps you make better use of your product.

CAUTION: A CAUTION indicates either potential damage to hardware or loss of data and tells you how to avoid the problem.

WARNING: A WARNING indicates a potential for property damage, personal injury, or death.

© 2020-2022 Dell Inc. or its subsidiaries. All rights reserved. Dell Technologies, Dell, and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be trademarks of their respective owners.



Chapter 1: Dell Technologies PowerEdge RAID Controller 11	
Features of PERC H755 adapter	9
Features of PERC H755 front SAS	9
Features of PERC H755N front NVMe	
Features of PERC H755 MX adapter	10
Features of PERC H750 adapter SAS	
Features of PERC H355 adapter SAS	
Features of PERC H355 front SAS	
Features of PERC H350 adapter SAS	
Features of PERC H350 Mini Monolithic SAS	
Operating systems supported by PERC 11 cards	
Technical specifications of PERC 11 cards	14
Thermal specifications	
	10
Comprehensive Embedded Management	18 1ຊ
Dell OpenManage Storage Management	
Human Interface Infrastructure Configuration Itility	
	10
Chapter 3: Features of PowerEdge RAID Controller 11	
Controller features	20
Non-Volatile Memory Express	20
Opal Security Management	
Hardware Root of Trust	
1 MB I/O	21
Autoconfigure RAID 0	
Disk roaming	
FastPath	
Non-RAID disks	23
Physical disk power management	23
Profile Management	23
Secure firmware update	
Snapdump	23
Virtual disk features	23
Virtual disk write cache policy	24
Virtual disk read cache policy	24
Virtual disk migration	25
Virtual disk initialization	
Full initialization	25
Fast initialization	25
Reconfigure virtual disks	
Background operations	
Background initialization	

Contents 3



Consistency checks	
Hard drive features	28
Self-Encrypting Disks	
Instant secure erase	
4 KB sector disk drives	
Fault tolerance	
The SMART feature	29
Patrol Read	
Physical disk failure detection	
Controller cache	
Battery Transparent Learn Cycle	32
Linux operating system device enumeration	
Chapter 4: Install and remove a PERC 11 card	
• Safety instructions	
Before working inside your system	
Remove the PERC H755 adapter	
Install the PERC H755 adapter	
Remove the PERC H755 front SAS card	
Install the PERC H755 front SAS card	38
Remove the PERC H755N front NVMe card	
Install the PERC H755N front NVMe card	41
Remove the PERC H755 MX adapter	42
Install the PERC H755 MX adapter	
Remove the PERC H750 adapter SAS	
Install the PERC H750 adapter SAS	45
Remove the PERC H355 adapter SAS	
Install the PERC H355 adapter SAS	47
Remove the PERC H355 front SAS	
Install the PERC H355 front SAS card	50
Remove the PERC H350 adapter SAS	51
Install the PERC H350 adapter SAS	52
Remove PERC H350 Mini Monolithic SAS	53
Install PERC H350 Mini Monolithic SAS	55
Chapter 5: Driver support for PERC 11	57
Creating the device driver media	57
Download and save PERC 11 drivers from the support site	57
Download and save PERC 11 drivers from the Dell Systems Service and Diagnostic Tools	
Windows driver installation	
Install PERC 11 driver while newly installing the Windows Server 2016 and later	
Install PERC 11 driver on which the Windows Server 2016 is already installed and later	
Update PERC 11 driver that runs on Windows Server 2016 or later	
Linux uriver installation	
Install of update a REIVI driver package using the KNOD support	
Install of update a KHIVI driver package using the KIVIH support	bU
Load the driver write installing at operating system	
Chapter 6: Firmware	62

4 Contents



Upgrade firmware controller using Dell Update Package (DUP).....

Chapter 7: Manage PERC 11 controllers using HII configuration utility	63
Enter the PERC 11 HII configuration utility	63
Exit the PERC 11 HII configuration utility	
Navigate to Dell PERC 11 configuration utility	64
View the HII Configuration utility dashboard	64
Configuration management	65
Auto Configure RAID 0	65
Create virtual disks	65
Create profile based virtual disk	
View disk group properties	67
Convert to Non-RAID disk	67
Delete configurations	67
Controller management	
Clear controller events	68
Save controller events	
Save debug log	68
Enable security	68
Disable security	
Change security settings	
Restore factory default settings	69
Auto configure behavior	
Manage controller profile	
Advanced controller properties	70
Virtual disk management	73
Virtual disk numbering	
Configure Virtual Disks	75
Perform expand virtual disk operation	75
Perform consistency check	76
Physical disk management	
View physical disk properties	76
Cryptographic erase	77
Physical disk erase	78
Assigning a global hot spare	78
Assigning a dedicated hot spare	
Convert to Non-RAID disk	79
Hardware components	
View battery properties	
View physical disks associated with an enclosure	
Security key management in HII configuration utility	81
Chapter 8: Security key and RAID management	82
Security key implementation	82 82
Local Key Management	
Create a security key	
Change Security Settings	
Disable security key	
Create a secured virtual disk	



Secure a pre-existing virtual disk Import a secured non-RAID disk Import a secured virtual disk Dell Technologies OpenManage Secure Enterprise Key Manager	
Import a secured non-RAID disk Import a secured virtual disk Dell Technologies OpenManage Secure Enterprise Key Manager	
Import a secured virtual disk Dell Technologies OpenManage Secure Enterprise Key Manager	
Dell Technologies OpenManage Secure Enterprise Key Manager	
	85 86 86
Supported controllers for OpenManage Secure Enterprise Key Manager	
Manage enterprise key manager mode	86
Disable enterprise key manager mode	
Manage virtual disks in enterprise key manager mode	
Manage non–RAID disks in enterprise key manager mode	
Transition of drives from local key management to enterprise key management (without supported firmware for PERC and iDRAC)	
Migrate of drives from local key management to enterprise key management (with supported firmware for PERC and iDRAC)	87
Chapter 9: Troubleshooting issues in PERC11 cards	88
Single virtual disk performance or latency in hypervisor configurations	
Configured disks removed or not accessible error message	
Dirty cache data error message	
Discovery error message	
Drive Configuration Changes Error Message	
Windows operating system installation errors	
Firmware fault state error message	
Foreign configuration found error message	
Foreign configuration not found in HII	
Degraded state of virtual disks	90
Memory errors	91
Preserved Cache State	
Security key errors	91
Secured foreign import errors	91
Failure to select or configure non Self-Encrypting Disks non-SED	
Failure to delete security key	92
Failure of Cryptographic Erase on encryption-capable physical disks	
General issues	
PERC card has yellow bang in Windows operating system device manager	92
PERC card not seen in operating systems	92
Issues in controller, battery, and disk when operating at low temperature	92
Physical disk issues	92
Physical disk in failed state	
Unable to rebuild a fault tolerant virtual disk	93
Fatal error or data corruption reported	
Multiple disks are inaccessible	93
Rebuilding data for a failed physical disk	93
Virtual disk fails during rebuild using a global hot spare	
Dedicated hot spare disk fails during rebuild	94
Redundant virtual disk fails during reconstruction	94
Virtual disk fails rebuild using a dedicated hot spare	94
Physical disk takes a long time to rebuild	94
Drive removal and insertion in the same slot generates a foreign configuration event	94
SMARI errors	

6 Contents



Smart error detected on a non-RAID disk	
Smart error detected on a physical disk in a non-redundant virtual disk	
Smart error detected on a physical disk in a redundant virtual disk	
Replace member errors	
Source disk fails during replace member operation	
Target disk fails during replace member operation	
A member disk failure is reported in the virtual disk which undergoes replace memb	per operation96
Linux operating system errors	
Virtual disk policy is assumed as write-through	
Unable to register SCSI device error message	
Drive indicator codes	
HII error messages	
Unhealthy Status of the drivers	
Rebuilding a drive during full initialization	
System reports more drive slots than what is available	
World Wide Number on drive sticker is not the same in applications	
Backplane firmware revision not changing in PERC interfaces after an update	
Chapter 10: Appendix RAID description	100
Summary of RAID levels	
RAID 10 configuration	
RAID terminology	
Disk striping	
Disk mirroring	
Spanned RAID levels	
Parity data	
Chapter 11: Getting help	
Recycling or End-of-Life service information	
Contacting Dell	
Locating the Express Service Code and Service Tag	
Receiving automated support with SupportAssist	
	40.0
Chapter 12: Documentation resources	



Dell Technologies PowerEdge RAID Controller 11

Dell Technologies PowerEdge RAID Controller 11, or PERC 11 is a series of RAID disk array controllers made by Dell for its PowerEdge servers. The PERC 11 series consists of the PERC H755 adapter, PERC H755 front SAS, PERC H755N front NVMe, PERC H750 adapter SAS, PERC H755 MX adapter, PERC H355 adapter SAS, PERC H355 front SAS, PERC H350 adapter SAS, and PERC H350 Mini Monolithic SAS cards that have the following characteristics:

- Provides reliability, high performance, and fault-tolerant disk subsystem management
- Offers RAID control capabilities including support for RAID levels 0, 1, 5, 6, 10, 50, 60
- Complies with Serial Attached SCSI (SAS) 3.0 providing up to 12 Gb/sec throughput
- Supports Dell-qualified Serial Attached SCSI (SAS), SATA hard drives, solid state drives (SSDs), and PCIe SSD (NVMe)
- Supports drive speeds of 8 GT/s and 16 GT/s at maximum x2 lane width for NVMe drives.
- **NOTE:** Mixing disks of different speeds (7,200 RPM, 10,000 RPM, or 15,000 RPM) and bandwidth (3 Gbps, 6 Gbps, or 12 Gbps) while maintaining the same drive type (SAS or SATA) and technology (hard drive or SSD) is supported.
- **NOTE:** Mixing NVMe drives with SAS and SATA is not supported. Also, mixing hard drives and SSDs in a virtual disk is not supported.
- **NOTE:** PERC H750 adapter SAS, PERC H355 front SAS, PERC H355 adapter SAS, PERC H350 adapter SAS, and PERC H350 Mini Monolithic SAS do not support NVMe drives.
- **NOTE:** RAID levels 5, 6, 50, and 60 are not supported on PERC H355 adapter SAS, PERC H355 front SAS, PERC H350 adapter SAS, and PERC H350 Mini Monolithic SAS.
- **NOTE:** PERC H350 Mini Monolithic SAS has form factor variations (Low Profile) for specific platforms. For more information, see your platform manuals.
- () NOTE: For the safety, regulatory, and ergonomic information that is associated with these devices, and for more information about the Integrated Dell Remote Access Controller (iDRAC) or Lifecycle Controller (LC) remote management, see your platform documentation.

Topics:

- Features of PERC H755 adapter
- Features of PERC H755 front SAS
- Features of PERC H755N front NVMe
- Features of PERC H755 MX adapter
- Features of PERC H750 adapter SAS
- Features of PERC H355 adapter SAS
- Features of PERC H355 front SAS
- Features of PERC H350 adapter SAS
- Features of PERC H350 Mini Monolithic SAS
- Operating systems supported by PERC 11 cards
- Technical specifications of PERC 11 cards
- Thermal specifications



Features of PERC H755 adapter

This section describes the features of PERC H755 adapter.



Figure 1. Features of PERC H755 adapter

- 1. Heatsink
- 3. Battery
- 5. Backplane connector B

- 2. PCle connector
- 4. Backplane connector A
- 6. Battery cable connector

Features of PERC H755 front SAS

This section describes the features of PERC H755 front SAS.



Figure 2. Features of PERC H755 front SAS

1. Battery

2. Backplane connector A

Dell Technologies PowerEdge RAID Controller 11 9

Fis.1416 Mov. 33

- 3. Power card edge connector
- 5. PCle input connector
- 7. Battery cable connector

- 4. Heatsink
- 6. Backplane connector B

Features of PERC H755N front NVMe



Figure 3. Features of PERC H755N front NVMe

- 1. Battery
- 3. Power card edge connector
- 5. Backplane connector A
- 7. Battery cable connector

- PCIe cable connector
 Heatsink
- 6. Backplane connector B
- Features of PERC H755 MX adapter



Figure 4. Features of PERC H755 MX adapter

- 1. Battery under cover
- 3. PCIe cable connector
- 5. Backplane connector B

- 2. Heatsink
- 4. Backplane connector A

10 Dell Technologies PowerEdge RAID Controller 11

Inserido ao protocolo **22.951.206-4** por: **Pedro Henrique Golin Linhares** em: 31/03/2025 12:49. A autenticidade deste documento pode ser validada no endereço: https://www.eprotocolo.pr.gov.br/spiweb/validarDocumento com o código: fc1c44dfe0686d5547ffd691199be699.



Features of PERC H750 adapter SAS



Figure 5. Features of PERC H750 adapter SAS

- 1. Heat sink
- 3. Battery cable connector

Battery
 Backplane connector A

5. PCle connector

Features of PERC H355 adapter SAS



Figure 6. Features of PERC H355 adapter SAS

- 1. Heat sink
- 3. Backplane connector A

Backplane connector B
 PCle connector

Dell Technologies PowerEdge RAID Controller 11 11



Features of PERC H355 front SAS



Figure 7. Features of H355 front SAS

- 1. PCle input connector
- 3. Backplane connector B
- 5. Power card edge connector

- 2. Heat sink
- 4. Backplane connector A

Features of PERC H350 adapter SAS



Figure 8. PERC H350 adapter SAS

- 1. Heat sink
- 2. Backplane connector A
- 3. PCle connector



Features of PERC H350 Mini Monolithic SAS



Figure 9. PERC H350 Mini Monolithic SAS

- 1. SAS cable connection
- 2. Heat sink

Operating systems supported by PERC 11 cards

See Dell Technologies Enterprise operating systems support for a list of supported operating systems by a specific server for the PERC 11 cards.

() NOTE: For the latest list of supported operating systems and driver installation instructions, see the operating system documentation at Operating Systems Documentation. For specific operating system service pack requirements, see the Drivers and Downloads section on the support site.



Technical specifications of PERC 11 cards

The following table lists the specifications of PERC 11 cards:

Table 1. Technical specifications of PERC 11 cards

Feature	PERC H755 adapter	PERC H755 front SAS	PERC H755N front NVMe	PERC H755 MX adapter	PERC H750 adapter SAS
RAID levels	0, 1, 5, 6, 10, 50, 60	0, 1, 5, 6, 10, 50, 60	0, 1, 5, 6, 10, 50, 60	0, 1, 5, 6, 10, 50 ,60	0, 1, 5, 6, 10, 50 ,60
Non-RAID	Yes	Yes	Yes	Yes	Yes
Enclosures per port	Not applicable				
Processor	Broadcom RAID- on-chip, SAS3916 chipset	Broadcom RAID- on-chip, SAS3916 chipset	Broadcom RAID-on- chip, SAS3916 chipset	Broadcom RAID- on-chip, SAS3916 chipset	Broadcom RAID-on- chip, SAS3916 chipset
Battery backup unit	Yes	Yes	Yes	Yes	Yes
Local Key Management security	Yes	Yes	Yes	Yes	Yes
Controller queue depth	5120	5120	5120	5120	5120
Secure enterprise key manager security	Yes	Yes	Yes	No	Yes
Non-volatile cache	Yes	Yes	Yes	Yes	Yes
Cache memory	8 GB DDR4 2666 MT/s cache	8 GB DDR4 2666 MT/s cache	8 GB DDR4 2666 MT/s cache	8 GB DDR4 2666 MT/s cache	8 GB DDR4 2666 MT/s cache
Cache function	Write back, write through, no read ahead, and read ahead	Write back, write through, no read ahead, and read ahead	Write back, write through, no read ahead, and read ahead	Write back, write through, no read ahead, and read ahead	Write back, write through, no read ahead, and read ahead
Max no of VDs in RAID mode	240	240	240	240	240
Max no of disk groups	240	240	240	240	240
Max no of VDs per disk group	16	16	16	16	16
Hot swap devices supported	Yes	Yes	Yes	Yes	Yes
Autoconfig	Yes	Yes	Yes	Yes	Yes
Hardware XOR engine	Yes	Yes	Yes	Yes	Yes
Online capacity expansion	Yes	Yes	Yes	Yes	Yes
Dedicated and global hot spare	Yes	Yes	Yes	Yes	Yes
Drives types	3 Gbps SATA, 6 Gbps SATA/SAS, and 12 Gbps SAS,	3 Gbps SATA, 6 Gbps SATA/SAS, and 12 Gbps SAS	Gen3 (8 GT/s) and Gen4 (16 GT/s) NVMe	3 Gbps SATA, 6 Gbps SATA/ SAS, and 12	3 Gbps SATA, 6 Gbps SATA/SAS, and 12 Gbps SAS

14 Dell Technologies PowerEdge RAID Controller 11

Inserido ao protocolo **22.951.206-4** por: **Pedro Henrique Golin Linhares** em: 31/03/2025 12:49. A autenticidade deste documento pode ser validada no endereço: https://www.eprotocolo.pr.gov.br/spiweb/validarDocumento com o código: fc1c44dfe0686d5547ffd691199be699.



Table 1. Technical specifications of PERC 11 cards (continued)

Feature	PERC H755 adapter	PERC H755 front SAS	PERC H755N front NVMe	PERC H755 MX adapter	PERC H750 adapter SAS
	Gen3 (8 GT/s), and Gen4 (16 GT/s) NVMe			Gbps SAS, Gen3 (8 GT/s), and Gen4 (16 GT/s) NVMe	
VD strip size	64 KB, 128 KB, 256 KB, 512 KB, 1 MB	64 KB, 128 KB, 256 KB, 512 KB, 1 MB	64 KB, 128 KB, 256 KB, 512 KB, 1 MB	64 KB, 128 KB, 256 KB, 512 KB, and 1 MB	64 KB, 128 KB, 256 KB, 512 KB, 1 MB
PCIe support	Gen 4	Gen 4	Gen 4	Gen 4	Gen 4
SAS/SATA maximum drive support	 Without SAS Expander: 16 drives per controller With SAS Expander: Limited by platform offerings 	 Without SAS Expander: 16 drives per controller With SAS Expander: Limited by platform offerings 	Not applicable	Limited by platform: 8 drives per controller	 Without SAS Expander: 8 drives per controller With SAS Expander: Limited by platform offerings
NVMe maximum drive support	 Without PCle Switch Expander: 8 drives per controller With PCle Switch Expander: Limited by platform offerings 	Not applicable	 Without PCle Switch Expander: 8 drives per controller With PCle Switch Expander: Limited by platform offerings 	Limited by platform:8 drives per controller	Not applicable

NOTE: PERC H755 adapter and PERC H755 MX supports either SAS, SATA, or NVMe drives depending on the backplane/ server configuration.

NOTE: PERC controller supports only conventional magnetic recording (CMR) drives, and does not support shingled magnetic recording (SMR) drives.

(i) NOTE: PERC H755 family of controllers currently support SEKM starting with firmware version 52.14.0-3901.

(i) NOTE: For information about the number of drives in a disk group per virtual disk, see Summary of RAID levels

(i) NOTE: As 14G PowerEdge Servers do not support Gen 4 speeds, PERC H750 adapter SAS will downtrain to Gen 3 speeds.

Table 2. Technical specifications of PERC 11 cards

Feature	PERC H355 adapter SAS	PERC H355 front SAS	PERC H350 adapter SAS	PERC H350 Mini Monolithic SAS
RAID levels	0, 1, 10	0, 1, 10	0, 1, 10	0, 1, 10
Non-RAID	Yes	Yes	Yes	Yes
Enclosures per port	Not applicable	Not applicable	Not applicable	Not applicable
Processor	Broadcom RAID- onchip, SAS3816 chipset	Broadcom RAID- onchip, SAS3816 chipset	Broadcom RAID- onchip, SAS3816 chipset	Broadcom RAID-onchip, SAS3816 chipset
Battery backup unit	No	No	No	No



Table 2. Technical specifications of PERC 11 cards (continued)

Feature	PERC H355 adapter SAS	PERC H355 front SAS	PERC H350 adapter SAS	PERC H350 Mini Monolithic SAS
Local Key Management security	No	No	No	No
Controller queue depth	1536	1536	1536	1536
Secure enterprise key manager security	No	No	No	No
Non-volatile cache	No	No	No	No
Cache memory	Not applicable	Not applicable	Not applicable	Not applicable
Cache function	Write through, no read ahead	Write through, no read ahead	write through, no read ahead	write through, no read ahead
Max no of VDs in RAID mode	32	32	32	32
Max no of disk groups	32	32	32	32
Max no of VDs per disk group	16	16	16	16
Hot swap devices supported	Yes	Yes	Yes	Yes
Autoconfig	Yes	Yes	Yes	Yes
Hardware XOR engine	Yes	Yes	Yes	Yes
Online capacity expansion	Yes	Yes	Yes	Yes
Dedicated and global hot spare	Yes	Yes	Yes	Yes
Drives types	3 Gbps SATA, 6 Gbps SATA/SAS, and 12 Gbps SAS, Gen3 (8 GT/s)	3 Gbps SATA, 6 Gbps SATA/SAS, and 12 Gbps SAS, Gen3 (8 GT/s)	3 Gbps SATA, 6 Gbps SATA/SAS, and 12 Gbps SAS, Gen3 (8 GT/s)	3 Gbps SATA, 6 Gbps SATA/SAS, and 12 Gbps SAS, Gen3 (8 GT/s)
VD strip size	64 KB	64 KB	64 KB	64 KB
PCle support	Gen 4	Gen 4	Gen 4	Gen 4
SAS/SATA maximum drive support	 Without SAS Expander: 16 With SAS Expander: Limited by platform offering 	 Without SAS Expander: 16 With SAS Expander: Limited by platform offering 	 Without SAS Expander: 8 drives per controller With SAS Expander: Limited by platform offering 	 Without SAS Expander: 8 drives per controller With SAS Expander: Limited by platform offering
NVMe maximum drive support	Not applicable	Not applicable	Not applicable	Not applicable

NOTE: As 14G PowerEdge Servers do not support Gen 4 speeds, PERC H350 adapter SAS and PERC H350 Mini Monolithic SAS will down train to Gen 3 speeds.

Thermal specifications

PERC 11 Controllers have an operating temperature range of 0C to 55C. System ambient temperatures may be less than or greater than these values.

16 Dell Technologies PowerEdge RAID Controller 11

Inserido ao protocolo **22.951.206-4** por: **Pedro Henrique Golin Linhares** em: 31/03/2025 12:49. A autenticidade deste documento pode ser validada no endereço: https://www.eprotocolo.pr.gov.br/spiweb/validarDocumento com o código: fc1c44dfe0686d5547ffd691199be699.


() NOTE: PERC Controllers may raise erroneous Battery, Disk, and Controller temperature errors if the controller is operating below the operational temperature range.



Applications and User Interfaces supported by PERC 11

PERC 11 card Management applications include the Comprehensive Embedded Management (CEM), Dell OpenManage Storage Management, The Human Interface Infrastructure (HII) configuration utility, and The PERC Command Line Interface (CLI). They enable you to manage and configure the RAID system, create and manage multiple disk groups, control and monitor multiple RAID systems, and provide online maintenance.

Topics:

- Comprehensive Embedded Management
- Dell OpenManage Storage Management
- Human Interface Infrastructure Configuration Utility
- The PERC Command Line Interface

Comprehensive Embedded Management

Comprehensive Embedded Management (CEM) is a storage management solution for Dell systems that enables you to monitor the RAID and network controllers installed on the system using iDRAC without an operating system installed on the system.

Using CEM enables you to do the following:

- Monitor devices with and without an operating systems installed on the system
- Provide a specific location to access monitored data of the storage devices and network cards
- Allows controller configuration for all PERC 11 cards
- **NOTE:** If you boot the system to HII (F2) or Lifecycle Controller (F10), then you cannot view the PERC cards on the CEM UI. The PERC cards are displayed on the CEM UI only after the system boot is complete.

(i) NOTE: It is not recommended that you create more than 8 VDs simultaneously with CEM.

Dell OpenManage Storage Management

Dell OpenManage Storage Management is a storage management application for Dell systems that provides enhanced features for configuring locally attached RAID disk storage. The Dell OpenManage storage management application enables you to perform controller and enclosure functions for all supported RAID controllers and enclosures from a single graphical or Command Line Interface (CLI). The User Interface (UI) is wizard-driven with features for novice and advanced users, and detailed online help. Using the Dell OpenManage storage management application, you can protect your data by configuring data-redundancy, assigning hot spares, or rebuilding failed physical disks. The fully featured CLI, which is available on select operating systems, allows you to perform RAID management tasks either directly from the console or through scripting.

(i) NOTE: For more information, see the Dell OpenManage Storage Management User's Guide at OpenManage Manuals

Human Interface Infrastructure Configuration Utility

The Human Interface Infrastructure (HII) configuration utility is a storage management application integrated into the system BIOS <F2>. It is used to configure and manage your Dell PowerEdge RAID Controller (PERC) virtual disks, and physical disks. This utility is independent of the operating system.

(i) NOTE: The BIOS configuration utility <Ctrl> <R> is not supported on PERC 11 cards.

18 Applications and User Interfaces supported by PERC 11



The PERC Command Line Interface

The PERC Command Line Interface (CLI) is a storage management application. This utility allows you to set up, configure, and manage your Dell PowerEdge RAID Controller (PERC) by using the Command Line Interface (CLI).

(i) NOTE: For more information, see Dell PowerEdge RAID Controller CLI Reference Guide at Storage Controllers Manuals.

Applications and User Interfaces supported by PERC 11 19



Features of PowerEdge RAID Controller 11

Topics:

- Controller features
- Virtual disk features
- Virtual disk initialization
- Reconfigure virtual disks
- Background operations
- Hard drive features
- Fault tolerance

Controller features

This section lists the following controller features supported on Dell Technologies PowerEdge RAID Controller 11 cards in detail:

- Non-Volatile Memory Express
- Opal Security Management
- Hardware Root of Trust
- 1 MB I/O
- Auto Configure RAID 0
- Disk roaming
- FastPath
- Non–RAID disks
- Physical disk power management
- Profile Management
- Secure firmware update
- Snapdump

Non-Volatile Memory Express

Non-Volatile Memory Express (NVMe) is a standardized, high-performance host controller interface and a storage protocol for communicating with non-volatile memory storage devices over the peripheral component interconnect express (PCle) interface standard. The PERC 11 controller supports up to 8 direct-attach NVMe drives. The PERC 11 controller is a PCle endpoint to the host, a PowerEdge server, and configured as a PCle root complex for downstream PCle NVMe devices connected to the controller.

NOTE: The NVMe drive on the PERC 11 controller shows up as a SCSI disk in the operating system, and the NVMe command line interface will not work for the attached NVMe drives.

Conditions under which a PERC supports an NVMe drive

- In NVMe devices the namespace identifier (NSID) with ID 1, which is (NSID=1) must be present.
- In NVMe devices with multiple namespace(s), you can use the drive capacity of the namespace with NSID=1.
- The namespace with NSID=1 must be formatted without protection information and cannot have the metadata enabled.
- PERC supports 512-bytes or 4 KB sector disk drives for NVMe devices.

20 Features of PowerEdge RAID Controller 11



Drive repair for NVMe initialization failure

If an NVME drive fails to initialize, the drive that is connected to PERC can be corrected in HII. The NVME initialization errors in the drives are listed as correctable and non-correctable errors in HII.

Repair drives with correctable NVMe initialization errors

Repair the drives with correctable NVMe initialization errors in HII to enable the drives to work properly.

About this task

Repairs can lead to permanent data loss in drives. Also, certain types of repairs can take a long time.

Steps

- 1. Log in to HII.
- 2. Go to Main Menu > Hardware Components > Enclosure Management. The drives with correctable and non-correctable errors are listed.
- **3.** Select the drive and click **Repair**. If the repair is successful, the drive is listed under physical drives and removed from the correctable error list. If the drive has other correctable errors, the drive is listed again in the correctable errors list.
- 4. If the repair is not successful, click **Repair** again.

(i) NOTE: In case you want to stop the repair, stop the repair from the Ongoing repairs list.

If the error is still not resolved or if the drive has other non-correctable errors, the drive is moved to the non-correctable error list.

Opal Security Management

Opal Security Management of Opal SED drives requires security key management support. You can use the application software or The Integrated Dell Remote Access Controller (iDRAC) to generate the security key that is set in the Opal drives and used as an authentication key to lock and unlock the Opal drives.

Hardware Root of Trust

Hardware RoT (RoT) builds a chain of trust by authenticating all the firmware components prior to its execution, and it permits only the authenticated firmware to perform and be flashed. The controller boots from an internal boot ROM (IBR) that establishes the initial root of trust and this process authenticates and builds a chain of trust with succeeding software using this root of trust.

1 MB I/O

PERC 11 controllers support a 1 MB I/O feature; if the capacity of I/O frame is greater than 1 MB, the I/O frame is broken into smaller chunks.

Autoconfigure RAID 0

The Autoconfigure RAID 0 feature creates a single drive RAID 0 on each hard drive that is in the ready state. For more information, see Auto Configure RAID 0.

NOTE: The Autoconfigure RAID 0 feature is not supported on PERC H355 adapter SAS, PERC H355 front SAS, PERC H350 adapter SAS, and PERC H350 Mini Monolithic SAS.



Autoconfigure behavior

The autoconfigure behavior automatically configures unconfigured drives during reboot and hot insertion. Unconfigured drives are configured according to the settings; but the configured drives remain unaffected. PERC 11 supports **Off and Non-RAID** settings.

Table 3. Autoconfigure behavior settings

Settings	Description
Off	Autoconfigure behavior is turned off.
Non-RAID	Unconfigured drives are configured as non-RAID disk during boot or during hot insertion; all the configured drives remain unaffected.
Off to Non–RAID disk	Unconfigured drives are converted to non-RAID disks; all the configured drives remain unaffected.
Non-RAID disk to Off	Unconfigured drives remain unconfigured good; all the configured drives remain unaffected.

NOTE: PERC H355 front SAS, PERC H355 adapter SAS, PERC H350 adapter SAS, and PERC H350 Mini Monolithic SAS converts an unconfigured good drive to non-RAID only if the drive has never been used before by that specific PERC.

Disk roaming

Disk roaming is when a physical disk is moved from one cable connection or backplane slot to another on the same controller. The controller automatically recognizes the relocated physical disks and logically places them in the virtual disks that are part of the disk group. If the physical disk is configured as a non-RAID disk, then the relocated physical disk is recognized as a non-RAID disk by the controller.

CAUTION: It is recommended that you perform disk roaming when the system is turned off.

CAUTION: Do not attempt disk roaming during RAID level migration (RLM) or online capacity expansion (OCE). This causes loss of the virtual disk.

Using disk roaming

About this task

Perform the following steps to use disk roaming:

Steps

- 1. Turn off the power to the system, physical disks, enclosures, and system components.
- 2. Disconnect power cables from the system.
- 3. Move the physical disks to desired positions on the backplane or the enclosure.
- 4. Perform a safety check. Make sure the physical disks are inserted properly.
- **5.** Turn on the system.

Results

The controller detects the RAID configuration from the configuration data on the physical disks.

FastPath

FastPath is a feature that improves application performance by delivering high I/O per second (IOPs) for solid-state drives (SSDs). The PERC 11 series of cards support FastPath.

To enable FastPath on a virtual disk, the cache policies of the RAID controller must be set to write-through and no read ahead. This enables FastPath to use the proper data path through the controller based on command (read/write), I/O size, and RAID type. For optimal solid-state drive performance, create virtual disks with strip size of 64 KB.

22 Features of PowerEdge RAID Controller 11



Non-RAID disks

A non-RAID disk is a single disk to the host, and not a RAID volume. The only supported cache policy for non-RAID disks is Write-Through.

Physical disk power management

Physical disk power management is a power-saving feature of PERC 11 series cards. The feature allows disks to be spun down based on disk configuration and I/O activity. The feature is supported on all rotating SAS and SATA disks, and includes unconfigured and hot-spare disks. The physical disk power management feature is disabled by default. You can enable the feature in the Dell Open Manage Storage Management application or in the Human Interface Infrastructure (HII) configuration utility. For more information on HII configuration and physical disk power management, see Enabling physical disk power management. For more information on using the Dell Open Manage Storage Management application, see the Dell OpenManage documentation at OpenManage Manuals.

Profile Management

PERC 11 supports the PD240 and PD64 profiles. It defines controller queue depth and the maximum number of physical and virtual disks.

Table 4. Supported profile on PERC 11

Feature	PD240	PD64
Controller	PERC H755 front SAS, PERC H755 MX adapter, and PERC H750 adapter SAS	PERC H355 front SAS, PERC H355 adapter SAS, PERC H350 adapter SAS, and PERC H350 Mini Monolithic SAS
Maximum virtual disk supported	240	32
Controller queue depth	5120	1536

Secure firmware update

This feature provides a cryptographic method of updating the firmware using an RSA encryption-decryption algorithm.

Only Dell-certified firmware is supported on your PERC controller.

Snapdump

The Snapdump feature provides the Dell support team with the debug information which can help to find the cause of firmware failure. In the instance of firmware failures, the firmware collects the logs and information at the time of failure, which are stored in a compressed file called a snapdump.

Snapdumps are also generated manually to provide additional debug information. When a snapdump is generated, it is stored in the controller's cache memory. This means in the event of a power loss the controller will offload the snapdump as part of its cache preservation mechanism. Snapdumps are preserved by default through four reboots before its deleted.

To generate a snapdump, change the snapdump, delete a snapdump, and to download a stored snapdump settings, see Dell PowerEdge RAID Controller CLI Reference Guide at Storage Controllers Manuals.

Virtual disk features

This section lists the following virtual disk features supported on PERC 11 cards in detail:

- Virtual disk read cache policies
- Virtual disk write cache policies
- Virtual disk migration



- Virtual disk initialization
- Reconfiguration of virtual disks
- Background operations

Virtual disk write cache policy

The write cache policy of a virtual disk determines how the controller handles writes to the virtual disk.

Table 5. Write cache policies

Feature	Description
Write-back	The controller sends a data transfer completion signal to the host when the controller cache has received all the data in a transaction. The controller then writes the cached data to the storage device in the background.
Write-through	The controller sends a data transfer completion signal to the host system when the disk subsystem has received all the data in a transaction. INOTE: Certain data patterns and configurations perform better with a write-through cache policy.

() NOTE: All RAID volumes are presented as write-through to the operating system (Windows and Linux) independent of the actual write cache policy of the virtual disk. PERC cards manage the data in cache independently of the operating system or any applications.

NOTE: Use the Dell OpenManage storage management application or the HII Configuration Utility to view and manage virtual disk cache settings.

Conditions under which write-back is employed

Write-back caching is used under all conditions in which the battery is present and in good condition.

Conditions under which forced write-back with no battery is employed

CAUTION: It is recommended that you use a power backup system when forcing write-back to ensure there is no loss of data if the system suddenly loses power.

Write-back mode is available when you select force write-back with no battery. When forced write-back mode is selected, the virtual disk is in write-back mode even if the battery is not present.

Virtual disk read cache policy

The read policy of a virtual disk determines how the controller handles reads to that virtual disk.

Table 6. Read policies

Feature	Description
Read ahead	Allows the controller to read sequentially ahead of requested data and to store the additional data in cache memory, anticipating that the data is required soon. This speeds up reads for sequential data, but there is slight improvement when accessing random data.
No read ahead	Disables the read ahead capability.

(i) NOTE: Adaptive read ahead is no longer supported. Selecting adaptive read ahead is equivalent to selecting the read ahead option.

24 Features of PowerEdge RAID Controller 11



Virtual disk migration

The PERC 11 series supports migration of virtual disks from one controller to another without taking the target controller offline. The controller can import RAID virtual disks in optimal, degraded, or partially degraded states. You cannot import a virtual disk that is offline. When a controller detects a configured physical disk, it marks the physical disk as foreign, and generates an alert indicating that a foreign disk was detected.

Disk migration pointers:

- Supports migration of virtual disks from H740P, H745, H745P MX, and H840 to the PERC 11 series except for H345.
- Supports migration of volumes that are created within the PERC 11 series.
- Does not support migration from the PERC 11 series to PERC H345, H740P, H745, H745P MX, and H840.
- Does not support migration from PERC H330, H730, and H830 to the PERC 11 series.

(i) NOTE: The source controller must be offline before performing the disk migration.

(i) NOTE: Importing non–RAID drives and uneven span RAID 10 virtual disks from PERC 9 to PERC 11 is not supported.

(i) NOTE: Disks cannot be migrated to older generations of PERC cards.

- **NOTE:** Importing secured virtual disks is supported as long as the appropriate local key management (LKM) is supplied or configured.
- NOTE: Virtual disk migration from PERC H755 adapter, PERC H755 front SAS, PERC H755N front NVMe, PERC H750 adapter SAS, PERC H755 MX adapter to PERC H350 adapter SAS, PERC H350 Mini Monolithic SAS, PERC H355 front SAS, and PERC H355 adapter SAS is not supported.

CAUTION: Do not attempt disk migration during RLM or online capacity expansion (OCE), this causes loss of the virtual disk.

Virtual disk initialization

PERC 11 series controllers support two types of virtual disk initialization:

- Full initialization
- Fast initialization

CAUTION: Initializing virtual disks erases files and file systems while keeping the virtual disk configuration intact.

Full initialization

Performing a full initialization on a virtual disk overwrites all blocks and destroys any data that previously existed on the virtual disk. Full initialization of a virtual disk eliminates the need for the virtual disk to undergo a Background Initialization (BGI). Full initialization can be performed after the virtual disk is created.

You can start a full initialization on a virtual disk by using the Slow Initialize option in the Dell OpenManage storage management application. For more information on using the HII Configuration Utility to perform a full initialization, see Configure virtual disk parameters.

(i) NOTE: If the system reboots during a full initialization, the operation aborts and a BGI begins on the virtual disk.

Fast initialization

A fast initialization on a virtual disk overwrites the first and last 8 MB of the virtual disk, clearing any boot records or partition information. The operation takes only 2–3 seconds to complete, but it is followed by BGI, which takes a longer time to complete. To perform a fast initialization using the HII Configuration Utility, see Configure virtual disk parameters.

NOTE: During full or fast initialization, the host cannot access the virtual disk. As a result, if the host attempts to access the virtual disk while it is initializing, all I/O sent by the host will fail.



NOTE: When using iDRAC to create a virtual disk, the drive undergoes fast initialization. During this process all I/O requests to the drive will respond with a sense key of **"Not Ready"** and the I/O operation will fail. If the operating system attempts to read from the drive as soon as it discovers the drive, and while the fast initialization is still in process, then the I/O operation fails and the operating system reports an I/O error.

Reconfigure virtual disks

An online virtual disk can be reconfigured in ways that expands its capacity and changes its RAID level.

NOTE: Spanned virtual disks such as RAID 50 and 60 califier be reconfigured.

(i) NOTE: Reconfiguring virtual disks typically impacts disk performance until the reconfiguration operation is complete.

Online Capacity Expansion (OCE) can be done in following ways:

 If there is a single virtual disk in a disk group and free space is available, the capacity of a virtual disk can be expanded within that free space. If multiple virtual disks exist within a common disk group, the capacities of those virtual disks cannot be expanded.

NOTE: Online capacity expansion is allowed on a disk group with a single virtual disk that begins at the start of the physical disk. It is not allowed when there is a free space at the beginning of a disk.

- 2. Add additional physical disks to a virtual disk to increase its capacity.
- **3.** After replacing all array members with larger drives than the original members, use the PERC CLI utility to expand the existing virtual disk to a larger size using the expandarray parameter. For more information, see Dell PowerEdge RAID Controller Command Line Interface Reference Guide.

RAID level migration (RLM) refers to changing a virtual disk's RAID level. Both RLM and OCE can be done simultaneously so that a virtual disk can simultaneously have its RAID level that is changed and its capacity increased. When an RLM or an OCE operation is complete, a reboot is not required.

CAUTION: Do not attempt disk migration during RLM or OCE operations. This causes loss of the virtual disk.

- **NOTE:** If an RLM or an OCE operation is in progress, then an automatic drive rebuild or copyback operation will not start until the operation is complete.
- **NOTE:** If the controller already contains the maximum number of virtual disks, you cannot perform a RAID level migration or capacity expansion on any virtual disk.
- **NOTE:** The controller changes the write cache policy of all virtual disks to write-through until the RLM or OCE operation is complete.
- (i) NOTE: You cannot initiate an OCE or an RLM on any virtual disk on the controller where a virtual disk with an ID of 0 exists.

See the following table for a list of RLM or OCE options: The source RAID level column indicates the virtual disk RAID level before the RLM or OCE operation and the target RAID level column indicates the RAID level after the RLM or OCE operation.

Table 7. RAID level migration

Source RAID Level	Target RAID Level	Number of Physical Disks (Beginning)	Number of Physical Disks (End)	Capacity Expansio n Possible	Description
RAID 0	RAID 0	1 or more	2 or more	Yes	Increases capacity by adding disks.
RAID 0	RAID 1	1	2	Yes	Converts a non-redundant virtual disk into a mirrored virtual disk by adding one disk.
RAID 0	RAID 5	1 or more	3 or more	Yes	Adds distributed parity redundancy; at least one disk must be added.



Table 7. RAID level migration (continued)

Source RAID Level	Target RAID Level	Number of Physical Disks (Beginning)	Number of Physical Disks (End)	Capacity Expansio n Possible	Description
RAID 0	RAID 6	1 or more	4 or more	Yes	Adds dual distributed parity redundancy; at least two disks must be added.
RAID 1	RAID 0	2	2 or more	Yes	Removes redundancy while increasing capacity.
RAID 1	RAID 5	2	3 or more	Yes	Maintains redundancy while adding capacity.
RAID 1	RAID 6	2	4 or more	Yes	Adds dual distributed parity redundancy and adds capacity.
RAID 5	RAID 0	3 or more	2 or more	Yes	Converts to a non-redundant virtual disk and reclaims disk space that is used for distributed parity data; one disk can be removed.
RAID 5	RAID 5	3 or more	4 or more	Yes	Increases capacity by adding disks.
RAID 5	RAID 6	3 or more	4 or more	Yes	Adds dual distributed parity redundancy; at least one disk needs to be added.
RAID 6	RAID 0	4 or more	2 or more	Yes	Converts to a non-redundant virtual disk and reclaims disk space that is used for distributed parity data; two disks can be removed.
RAID 6	RAID 5	4 or more	3 or more	Yes	Removes one set of parity data and reclaims disk space used for it; one disk can be removed.
RAID 6	RAID 6	4 or more	5 or more	Yes	Increases capacity by adding disks.
RAID 10	RAID 10	4 or more	6 or more	Yes	Increases capacity by adding disks; an even number of disks must be added.

(i) NOTE: You cannot perform a RAID level migration and expansion on RAID levels 50 and 60.

Background operations

Background initialization

Background initialization (BGI) is an automated process that writes parity or mirror data on newly created virtual disks. BGI does not run on RAID 0 virtual disks. You can control the BGI rate in the Dell OpenManage storage management application. Any change to the BGI rate does not take effect until the next BGI is performed.

() NOTE:

- You cannot disable BGI permanently. If you cancel BGI, it automatically restarts within five minutes.
- Unlike full or fast initialization of virtual disks, background initialization does not clear data from the physical disks.
- Consistency Check (CC) and BGI typically cause some loss in performance until the operation completes.



 PERC H355 adapter SAS, PERC H355 front SAS, PERC H350 adapter SAS, and PERC H350 Mini Monolithic SAS background operations will not run until the operating system boots.

Consistency check and BGI perform similar functions in that they both correct parity errors. However, CC reports data inconsistencies through an event notification, while BGI does not. You can start CC manually, but not BGI.

Consistency checks

Consistency Check (CC) is a background operation that verifies and corrects the mirror or parity data for fault tolerant virtual disks. It is recommended that you periodically run a consistency check on virtual disks.

You can manually start a CC using the HII Configuration Utility or the Dell OpenManage storage management application. You can schedule a CC to run on virtual disks using the Dell OpenManage storage management application. To start a CC using the HII Configuration Utility, see Perform consistency check.

(i) NOTE: CC or BGI typically causes some loss in performance until the operation completes.

CC and BGI both correct parity errors. However, CC reports data inconsistencies through an event notification, while BGI does not. You can start CC manually, but not BGI.

Hard drive features

This section lists the following hard drive features supported on PERC 11 cards in detail:

- Self-Encrypting Disks (SED)
- Instant Secure Erase (ISE)
- 4 KB sector disk drives

Self-Encrypting Disks

Select PERC 11 cards support self-encrypting disks (SEDs) for protection of data against loss or theft of SEDs. For information about cards supported, see Technical specifications. Protection is achieved by the use of encryption technology on the drives. There is one security key per controller. You can manage the security key using local key management (LKM) or OpenManage Secure Enterprise Key Manager also referred as Secure Enterprise Key Manager (SEKM). The controller use the security key to lock and unlock access to encryption-capable physical disks. To take advantage of this feature, you must:

- Have SEDs in your system, and
- Create a security key.

PERC cannot use SEDs that are secured by a non-PERC entity. Ensure that the SED is reprovisioned in an applicable manner by that non-PERC entity before connecting to PERC.

For more information, see the Security key and RAID management section.

(i) NOTE: You cannot enable security on non-optimal virtual disks.

NOTE: PERC 11 supports Trusted Computing Group Enterprise (TCG) Security Subsystem Classes (SSC) SAS or SATA SED drives and TCG Opal SSC NVMe drives.

Instant secure erase

Instant Secure Erase (ISE) drives use the same encryption technology as SED drives but do not allow the encryption key to be secured. The encryption technology allows the drive to be re-purposed and securely erased using the cryptographic erase function.

(i) NOTE: ISE drives do not provide protection against theft.

4 KB sector disk drives

PERC 11 controllers support 4 KB sector disk drives, which enables you to efficiently use the storage space.

28 Features of PowerEdge RAID Controller 11



Before installing Windows on 4 KB sector disk drives, see Windows operating system installation errors.

NOTE: Mixing 512–byte native and 512–byte emulated drives in a virtual disk is allowed, but mixing 512–byte and 4 KB native drives in a virtual disk is not allowed.

(i) NOTE: 4 K is only supported in UEFI mode and not legacy BIOS.

(i) NOTE: 4 K devices do not appear under the select boot device option. For more information, see Enable boot support.

Fault tolerance

The PERC 11 series supports the following:

- Self-Monitoring and Reporting Technology (SMART)
- Patrol read
- Physical disk failure detection
- Controller cache
- Battery Transparent Learn Cycle

The next sections describe some methods to achieve fault tolerance.

The SMART feature

The SMART feature monitors certain physical aspects of all motors, heads, and physical disk electronics to help detect predictable hard drive failures. Data on SMART compliant hard drives can be monitored to identify changes in values and determine whether the values are within threshold limits. Many mechanical and electrical failures display some degradation in performance before failure.

A SMART failure is also referred to as predicted failure. There are numerous factors that are predicted physical disk failures, such as a bearing failure, a broken read/write head, and changes in spin-up rate. In addition, there are factors that are related to read/write surface failure, such as seek error rate and excessive bad sectors.

NOTE: For detailed information about SCSI interface specifications, see t10.org and for detailed information about SATA interface specifications, see t13.org.

NOTE: PERC H355 adapter SAS, PERC H355 front SAS, PERC H350 adapter SAS, and PERC H350 Mini Monolithic SAS controllers do not monitor predictive failures for non-RAID disks.

Automatic Replace Member with predicted failure

A replace member operation can occur when there is a SMART predictive failure reporting on a physical disk in a virtual disk. The automatic replace member is initiated when the first SMART error occurs on a physical disk that is part of a virtual disk. The target disk needs to be a hot spare that qualifies as a rebuild disk. The physical disk with the SMART error is marked as failed only after the successful completion of the replace member. This prevents the array from reaching degraded state.

If an automatic replace member occurs using a source disk that was originally a hot spare (that was used in a rebuild), and a new disk is added and set as a target disk for the replace member operation, the hot spare drive will revert to the hot spare state after the replace member operation successfully completes.

(i) NOTE: To enable automatic replace member, use the Dell storage management application.

Patrol Read

The Patrol read feature is designed as a preventative measure to ensure physical disk health and data integrity. Patrol read scans and resolves potential problems on configured physical disks. The Dell storage management applications can be used to start patrol read and change its behavior.

The following is an overview of patrol read behavior:

- Patrol read runs on all disks on the controller that are configured as part of a virtual disk, including hot spares.
- Patrol read does not run on physical disks that are not part of a virtual disk or are in Ready state.



- The amount of controller resources dedicated to patrol read operations adjusts based on the number of outstanding disk I/O
 operations. For example, if the system is processing a large number of I/O operations, then patrol read uses fewer resources
 to allow the I/O to take a higher priority.
- Patrol read does not run on disks that are involved in any of the following operations:
 - o Rebuild
 - Replace member
 - Full or background initialization
 - CC
 - RLM or OCE

(i) NOTE: By default, patrol read automatically runs every seven days on configured SAS and SATA hard drives.

For more information about patrol read, see the Dell OpenManage documentation at OpenManage Manuals.

Physical disk failure detection

If a disk fails and it is replaced with a new disk, the controller will automatically start a rebuild on the new disk. See, Configured slot behavior. Automatic rebuilds can also occur with hot spares. If you have configured hot spares, the controller will automatically try to use them to rebuild the degraded virtual disk.

Using persistent hot spare slots

(i) NOTE: The persistent hot spare slot feature is disabled by default.

The PERC 11 series can be configured so that the system backplane or storage enclosure disk slots are dedicated as hot spare slots. This feature can be enabled using the Dell storage management application.

Once enabled, any slots with hot spares configured automatically become persistent hot spare slots. If a hot spare disk fails or is removed, a replacement disk that is inserted into the same slot automatically becomes a hot spare with the same properties as the one it is replacing. If the replacement disk does not match the disk protocol and technology, it does not become a hot spare.

For more information on persistent hot spares, see the Dell OpenManage documentation at OpenManage Manuals.

Configured slot behavior

This feature is similar to persistent hot spare slot behavior. If a redundant VD is configured to the system and if a drive is replaced, the configured slot will automatically rebuild or copyback on the inserted drive regardless of the data on the drive. This operation will overwrite the data on the drive.

Table 8. Drive state/operation

Drive state/operation	Unconfigured slot	Slot configured in VD
Insert unconfigured drive into the system	Ready	Rebuild or copyback start
Insert configured drive into the system	Foreign	Rebuild or copyback startOriginal drive data lost
Insert configured locked drive into the system (unlockable)	Foreign	Cryptographic Erase (If configured VD is not secured) • Rebuild or copyback start • Original drive data lost
Insert locked drive into the system (non-unlockable)	Foreign locked	Foreign locked

Physical disk hot swapping

Hot swapping is the manual replacement of a disk while the PERC 11 series cards are online and performing their normal functions. The following requirements must be met before hot swapping a physical disk:

Inserido ao protocolo **22.951.206-4** por: **Pedro Henrique Golin Linhares** em: 31/03/2025 12:49. A autenticidade deste documento pode ser validada no endereço: https://www.eprotocolo.pr.gov.br/spiweb/validarDocumento com o código: fc1c44dfe0686d5547ffd691199be699.



- The system backplane or enclosure must support hot swapping for the PERC 11 series cards.
- The replacement disk must be of the same protocol and disk technology. For example, only a SAS hard drive can replace a SAS hard drive and only a NVMe drive can replace a NVMe drive.

Using replace member and revertible hot spares

The replace member functionality allows a previously commissioned hot spare to revert to a usable hot spare. When a disk failure occurs within a virtual disk, an assigned hot spare, dedicated, or global, is commissioned and begins rebuilding until the virtual disk is optimal. After the failed disk is replaced in the same slot and the rebuild to the hot spare is complete, the controller automatically starts to copy data from the commissioned hot spare to the newly inserted disk. After the data is copied, the new disk is a part of the virtual disk and the hot spare is reverted to being a ready hot spare. This allows hot spares to remain in specific enclosure slots. While the controller is reverting the hot spare, the virtual disk remains optimal. The controller automatically reverts a hot spare only if the failed disk is replaced with a new disk in the same slot. If the new disk is not placed in the same slot, a manual replace member operation can be used to revert a previously commissioned hot spare.

NOTE: A replace member operation typically causes a temporary impact to disk performance. Once the operation completes, performance returns to normal.

Controller cache

The PERC 11 series of cards contain local DRAM on the controllers. This DRAM can cache I/O operations for Write Back, Read Ahead virtual disks to improve the performance.

NOTE: Virtual disks consisting of SSDs may not see a difference in performance using controller cache and may benefit by Fastpath.

I/O workload that is slow to HDDs, such as random 512 B and 4 kB, may take some time to flush cached data. Cache is flushed periodically but for configuration changes or system shutdown, the cache is required to be flushed before the operation can be completed. It can take several minutes to flush cache for some workloads depending on the speed of the HDDs and the amount of data in the cache.

The following operations require a complete cache flush:

- Configuration changes (add or delete VDs, VD cache setting changes, foreign configuration scan, and import)
- System reboot or shutdown
- Abrupt power loss causing cache preservation
- () NOTE: The iDRAC or OpenManage periodically scans for the foreign configurations when the foreign disks are present. This action degrades the performance. If a foreign disk is present, it is recommended that you import, clear, or remove the foreign disk to prevent an impact on the performance.

Controller cache preservation

The controller is capable of preserving its cache in the event of a system power outage or improper system shutdown. The PERC 11 series controller is attached to a battery backup unit (BBU) that provides backup power during system power loss to preserve the controller's cache data.

Cache preservation with non-volatile cache

The non-volatile cache (NVC) allows controller cache data to be stored indefinitely. If the controller has data in the cache memory during a power outage or improper system shutdown, a small amount of power from the battery is used to transfer the cache data to non-volatile flash storage where it remains until power is restored and the system is booted. If the cache preservation process is interrupted by power-on, the controller may request an extra reset during the boot to complete the process. The system displays a message during boot as Dell PERC at Bus <X> Dev <Y> has requested a system reset. System will reboot in 5 seconds.



Recovering cache data

About this task

Complete these steps if a system power loss or improper system shutdown has occurred.

Steps

- 1. Restore the system power.
- **2.** Boot the system.
- **3.** When preserved cache exists on the controller, an error message is shown. For more information about how to recover cache, see Preserved Cache State.

Battery Transparent Learn Cycle

A transparent learn cycle is a periodic operation that calculates the charge that is remaining in the battery to ensure that there is sufficient energy. The operation runs automatically, and causes no impact to the system or controller performance.

The controller automatically performs the transparent learn cycle (TLC) on the battery to calibrate and gauge its charge capacity once every 90 days. The operation can be performed manually if required.

NOTE: Virtual disks stay in write-back mode, if enabled, during transparent learn cycle. When the TLC completes, the controller sets the next TLC to +90 days.

Transparent Learn Cycle completion time

The time frame for completion of a learn cycle is a function of the battery charge capacity and the discharge and charge currents used. Typical time completion for a transparent learn cycle is between 4 to 8 hours. If the learn cycle is interrupted mid cycle, it begins at a new cycle.

Conditions for replacing the battery

The PERC battery is marked failed when the state or health of the battery is declared bad. If the battery is declared failed, then all the virtual disks in write-back mode transitions to write-through mode, and the firmware runs learn cycles in subsequent reboots until the battery is replaced. On replacing the battery, virtual disk transitions to write-back mode.

Linux operating system device enumeration

Virtual disks and non-RAID disks are presented to the operating system as SCSI devices. The operating system enumerates these devices based on the SCSI target device ID.

Enumeration order for PERC H355 adapter SAS, PERC H355 front SAS, PERC H350 adapter SAS, and PERC H350 Mini Monolithic SAS

Steps

- 1. Non-RAID disks are enumerated first.
- Virtual disks (VDs) are enumerated second, based on virtual disk target ID. Target IDs are assigned to the VDs in the ascending order when they are created. The first created VD is assigned the lowest available target ID, and the last created VD is assigned the highest available target ID. The first created VD is discovered first by the operating system.
 NOTE: The PERC H355 adapter SAS, PERC H355 front SAS, PERC H350 adapter SAS, and PERC H350 Mini

Monolithic SAS non-RAID disks may not appear in the slot order.



Enumeration order for PERC H755 front SAS, PERC H755 front SAS, PERC H755N front NVMe, PERC H750 adapter SAS, and PERC H755 MX adapter

This section describes the order of enumerating PERC H-series controlers.

Steps

- 1. Non-RAID disks are enumerated first based on slot ID.
- 2. Virtual disks (VDs) are enumerated, second based on the virtual disk target ID. Target IDs are assigned to the VDs in the descending order when they are created. The first created VD is assigned the highest available target ID, and the last created VD is assigned the lowest available target ID. Therefore, the last created VD is discovered first by the operating system.
 - () **NOTE:** Operating system enumeration may not be in this order if virtual disks or non-RAID disks are created while the operating system is running. The operating system may name devices based on the order in which they were created resulting in the operating system enumeration changing after reboot. It is recommended to reboot the system for the final device enumeration after creating any virtual disks or non-RAID disks.



Install and remove a PERC 11 card

Topics:

- Safety instructions
- Before working inside your system
- Remove the PERC H755 adapter
- Install the PERC H755 adapter
- Remove the PERC H755 front SAS card
- Install the PERC H755 front SAS card
- Remove the PERC H755N front NVMe card
- Install the PERC H755N front NVMe card
- Remove the PERC H755 MX adapter
- Install the PERC H755 MX adapter
- Remove the PERC H750 adapter SAS
- Install the PERC H750 adapter SAS
- Remove the PERC H355 adapter SAS
- Install the PERC H355 adapter SAS
- Remove the PERC H355 front SAS
- Install the PERC H355 front SAS card
- Remove the PERC H350 adapter SAS
- Install the PERC H350 adapter SAS
- Remove PERC H350 Mini Monolithic SAS
- Install PERC H350 Mini Monolithic SAS

Safety instructions

\triangle	CAUTION: Ensure that two or more people lift the system horizontally from the box and place it on a flat surface, rack lift, or into the rails.
Δ	WARNING: Opening or removing the PowerEdge server cover while the server is powered on may expose you to a risk of electric shock.
Δ	WARNING: Do not operate the server without the cover for a duration exceeding five minutes. Operating the system without the system cover can result in component damage.
()	NOTE: Many repairs may only be done by a certified service technician. You should only perform troubleshooting and simple repairs as authorized in your product documentation, or as directed by the online or telephone service and support team. Damage due to servicing that is not authorized by Dell is not covered by your warranty. Read and follow the safety instructions that are shipped with your product.
\triangle	CAUTION: To ensure proper operation and cooling, all system bays and fans must always be populated with a component or a blank.
(j	NOTE: It is recommended that you always use an antistatic mat and antistatic strap while working on components inside the server.
(j	NOTE: To ensure proper operation and cooling, all system bays and fans must always be populated with a component or a blank.
(j	NOTE: While replacing the hot swappable PSU, after next server boot, the new PSU automatically updates to the same firmware and configuration of the replaced one.



Before working inside your system

Prerequisites

Follows the steps listed in Safety instructions.

Steps

- 1. Power off the system and all attached peripherals.
- 2. Disconnect the system from the electrical outlet and disconnect the peripherals.
- **3.** If applicable, remove the system from the rack.

For more information, see the Rail Installation Guide relevant to your rail solutions at PowerEdge Manuals.

4. Remove the system cover.

Remove the PERC H755 adapter

Prerequisites

CAUTION: Many repairs may only be done by a certified service technician. You should only perform troubleshooting and simple repairs as authorized in your product documentation, or as directed by the online or telephone service and support team. Damage due to servicing that is not authorized by Dell is not covered by your warranty. Read and follow the safety instructions that are shipped with your product.

NOTE: It is recommended that you always use a static mat and static strap while working on components in the interior of the system.

Steps

1. Turn off the system, including any attached peripherals, and disconnect the system from the electrical outlet and peripherals.

NOTE: Perform a graceful shutdown of the system to ensure data in the cache is moved to the disk before the controller is removed.

- 2. Open the system.
- 3. Locate the PERC card in the expansion riser on the system board.

CAUTION: To prevent damage to the card, you must hold the card by its edges only.

- 4. Unfasten and lift the riser from the system board. Remove the PERC card.
- 5. Disconnect any cables connected to the card:
 - a. Press down and hold the metal tab on the cable connector.
 - **b.** Pull the cable out of the connector.
- 6. Replace the storage controller card and reconnect the data cables before placing them in the riser. For more information on installing the card, see Install PERC H755 adapter.
- 7. Reinstall the riser on the system board and fasten the riser.
- 8. Close the system.
- 9. Reconnect the system to its electrical outlet and turn the system on, including any attached peripherals.





Figure 10. Remove the PERC H755 adapter

Install the PERC H755 adapter

Prerequisites

CAUTION: Many repairs may only be done by a certified service technician. You should only perform troubleshooting and simple repairs as authorized in your product documentation, or as directed by the online or telephone service and support team. Damage due to servicing that is not authorized by Dell is not covered by your warranty. Read and follow the safety instructions that are shipped with your product.

NOTE: It is recommended that you always use a static mat and static strap while working on components in the interior of the system.

Steps

- 1. Turn off the system, including any attached peripherals, and disconnect the system from the electrical outlet.
- 2. Open the system.
- **3.** Align the card-edge connector with the connector on the system board.

\triangle CAUTION: To prevent damage to the card, you must hold the card by its edges only.

- 4. Press the card-edge down until the card is fully seated.
- 5. Connect the data cable connectors to the card.
- 6. Route the data cable through the channel on the inner side of the chassis to the backplane.
- 7. Attach the connector to the corresponding connector on the backplane as labeled on the controller.
- 8. Close the system.
- 9. Reconnect the system to its electrical outlet and turn the system on, including any attached peripherals.





Figure 11. Install the PERC H755 adapter

Remove the PERC H755 front SAS card

Prerequisites

CAUTION: Many repairs may only be done by a certified service technician. You should only perform troubleshooting and simple repairs as authorized in your product documentation, or as directed by the online or telephone service and support team. Damage due to servicing that is not authorized by Dell is not covered by your warranty. Read and follow the safety instructions that are shipped with your product.

NOTE: It is recommended that you always use a static mat and static strap while working on components in the interior of the system.

Steps

1. Turn off the system, including any attached peripherals, and disconnect the system from the electrical outlet and peripherals.

NOTE: Perform a graceful shutdown of the system to ensure data in the cache is moved to the disk before the controller is removed.

- 2. Open the system.
- 3. Locate the PERC card in the controller carrier at the front of the system.

CAUTION: To prevent damage to the card, you must hold the card by its edges only.

4. Unscrew the fasteners on the controller carrier and slide the carrier away from the backplane, disconnecting the controller from the backplane.

If you are removing a PERC H755 front SAS controller in the upside down orientation, you must remove both the backplane and the controller at the same time because of the limited clearance available:

- a. Uninstall all drives from the backplane.
- b. Disconnect all cables between the PERC and the backplane.
- c. Lift the backplane and PERC from the system.
- **5.** Disconnect any cables connected to the card:
 - a. Press down and hold the metal tab on the cable connector.
 - **b.** Pull the cables out of the connector.
- 6. Remove the PERC controller from the controller carrier.

Install and remove a PERC 11 card 37



- 7. Insert the replacement controller into the carrier and secure it with the appropriate screws.
- 8. Take the replacement storage controller and reconnect the cables before reconnecting it to the backplane. If you are removing a PERC H755 front SAS controller in the upside down orientation, reattach the PERC controller to the backplane first before reinstalling the backplane into the system. For more information on installing the card, see Install PERC H755 front SAS card.
- 9. Close the system.
- 10. Reconnect the system to its electrical outlet and turn the system on, including any attached peripherals.



Figure 12. Remove the PERC H755 front SAS card

Install the PERC H755 front SAS card

Prerequisites

- CAUTION: Many repairs may only be done by a certified service technician. You should only perform troubleshooting and simple repairs as authorized in your product documentation, or as directed by the online or telephone service and support team. Damage due to servicing that is not authorized by Dell is not covered by your warranty. Read and follow the safety instructions that are shipped with your product.
- **NOTE:** It is recommended that you always use a static mat and static strap while working on components in the interior of the system.

Steps

1. Turn off the system, including any attached peripherals, and disconnect the system from the electrical outlet.



NOTE: Perform a graceful shutdown of the sled to ensure that data in the cache is moved to the disk before the controller is removed.

- 2. Open the system.
- 3. Connect the PERC card to the carrier and ensure that the screws are properly fastened in place.

CAUTION: To prevent damage to the card, hold the card by its edges only.

- **4.** Align the carrier with the guide pins until the controller is securely seated.
- 5. Slide the card into the connector until it is fully seated in the connector. Tighten the screws on the carrier that connect to the chassis to secure the carrier.
- 6. Connect the cable connectors to the card.

NOTE: Ensure that you connect the cable according to the connector labels on the cable. The cable does not function properly if reversed.

- 7. Close the system.
- 8. Reconnect the system to its electrical outlet and turn on the system and any attached peripherals.



Figure 13. Install the PERC H755 front SAS card

Remove the PERC H755N front NVMe card

Prerequisites

CAUTION: Many repairs may only be done by a certified service technician. You should only perform troubleshooting and simple repairs as authorized in your product documentation, or as directed by the online or

Install and remove a PERC 11 card 39



telephone service and support team. Damage due to servicing that is not authorized by Dell is not covered by your warranty. Read and follow the safety instructions that are shipped with your product.

NOTE: It is recommended that you always use a static mat and static strap while working on components in the interior of the system.

Steps

- 1. Turn off the system, including any attached peripherals, and disconnect the system from the electrical outlet and peripherals.
 - **NOTE:** Perform a graceful shutdown of the system to ensure that data in the cache is moved to the disk before the controller is removed.
- 2. Open the system.
- 3. Locate the PERC card in the controller carrier at the front of the system.

CAUTION: To prevent damage to the card, you must hold the card by its edges only.

4. Unscrew the fasteners on the controller carrier, and slide the carrier away from the backplane to disconnect the controller from the backplane.

If you are removing a PERC H755N front NVMe controller in the upside down orientation, you must remove both the backplane and the controller at the same time because of the limited clearance available:

- **a.** Uninstall all drives from the backplane.
- **b.** Disconnect all cables between the PERC and the backplane.
- c. Lift the backplane and PERC from the system.
- 5. Disconnect any cables connected to the card:
 - a. Press down and hold the metal tab on the cable connector.
 - **b.** Pull the cable out of the connector.
- 6. Remove the PERC controller from the controller carrier.
- 7. Insert the replacement controller into the carrier and secure it with the appropriate screws.
- 8. Take the replacement storage controller and reconnect the cable before reconnecting it to the backplane.

If you are removing a PERC H755 front NVMe controller in the upside down orientation, reattach the PERC controller to the backplane first before reinstalling the backplane into the system.

- 9. Close the system.
- 10. Reconnect the system to its electrical outlet and turn the system on, including any attached peripherals.





Figure 14. Remove the PERC H755N front NVMe card

Install the PERC H755N front NVMe card

Prerequisites

- CAUTION: Many repairs may only be done by a certified service technician. You should only perform troubleshooting and simple repairs as authorized in your product documentation, or as directed by the online or telephone service and support team. Damage due to servicing that is not authorized by Dell is not covered by your warranty. Read and follow the safety instructions that are shipped with your product.
- **NOTE:** It is recommended that you always use a static mat and static strap while working on components in the interior of the system.

Steps

- 1. Turn off the system, including any attached peripherals, and disconnect the system from the electrical outlet.
 - () NOTE: Perform a graceful shutdown of the sled to ensure that data in the cache is moved to the disk before the controller is removed.
- 2. Open the system.
- **3.** Connect the PERC card to the carrier and ensure the screws are properly fastened in place.

CAUTION: To prevent damage to the card, hold the card by its edges only.

4. Align the carrier with the guide pins until the controller is securely seated.

Install and remove a PERC 11 card 41



- 5. Slide the card until it is fully seated in the connector. Tighten the screws on the carrier that connect to the chassis to secure the carrier.
- ${\bf 6.}\$ Connect the cable connectors to the card.

() NOTE: Ensure that you connect the cable according to the connector labels on the cable. The cable does not function properly if reversed.

- 7. Close the system.
- 8. Reconnect the system to its electrical outlet and turn on the system and any attached peripherals.



Figure 15. Install the PERC H755N front NVMe card

Remove the PERC H755 MX adapter

Prerequisites

CAUTION: Many repairs may only be done by a certified service technician. You should only perform troubleshooting and simple repairs as authorized in your product documentation, or as directed by the online or telephone service and support team. Damage due to servicing that is not authorized by Dell is not covered by your warranty. Read and follow the safety instructions that are shipped with your product.

CAUTION: To prevent damage to the card, hold the card by its edges only.

() NOTE: It is recommended that you always use a static mat and static strap while working on components in the interior of the system.



Steps

- 1. Turn off the sled, including any attached peripherals, and remove the sled from the MX chassis.
 - **NOTE:** Perform a graceful shutdown of the system to ensure that data in the cache is moved to the disk before the controller is removed.
- 2. Open the sled.
- **3.** Locate the PERC card on the system board.

CAUTION: To prevent damage to the card, hold the card by its edges only.

- 4. Using the blue tab, rotate the lever of the controller.
- 5. Pull the release lever upward to disengage the controller from the connector.
- 6. Disconnect the cable from the card. To disconnect the cable:
 - **a.** Press and hold the metal tab on the cable connector.
 - **b.** Pull the cable out of the connector.
- 7. Lift the card from the system board.
- Replace the storage controller card and connect the cable. For information on installing the card, see Install the PERC H755 MX adapter.
- 9. Close the sled.
- 10. Insert the sled into the MX chassis and turn on the system and any attached MX chassis peripherals.



Figure 16. Remove the PERC H755 MX adapter

Install the PERC H755 MX adapter

Prerequisites

CAUTION: Many repairs may only be done by a certified service technician. You should only perform troubleshooting and simple repairs as authorized in your product documentation, or as directed by the online or

Install and remove a PERC 11 card 43



telephone service and support team. Damage due to servicing that is not authorized by Dell is not covered by your warranty. Read and follow the safety instructions that are shipped with your product.

NOTE: It is recommended that you always use a static mat and static strap while working on components in the interior of the system.

Steps

- 1. Turn off the sled and any attached peripherals, and remove the sled from the MX chassis.
- 2. Open the sled.
- **3.** Connect the backplane data cable connector to the card.
 - **NOTE:** Ensure that you connect the cable according to the connector labels on the cable. The cable does not function properly if reversed.
- **4.** Align the bracket notches with the tabs on the sides of the sled chassis and align the PERC card connector with the connector on the system board.

CAUTION: To prevent damage to the card, hold the card by its edges only.

- 5. Press the PERC card into the connector until it is firmly seated.
- 6. Press the release lever to secure the card to the sled.

(i) NOTE: The pin on the release lever secures the card to the chassis of the sled.

- 7. Route the data cable through the clip on the card and through the channel on the inner side of the chassis.
- 8. Attach the connector to the corresponding connector on the backplane as labeled in the controller.
- 9. Close the sled.
- 10. Insert the sled into the MX chassis and turn on the system and any attached MX chassis peripherals.



Figure 17. Install the PERC H755 MX adapter



Remove the PERC H750 adapter SAS

Prerequisites

CAUTION: Many repairs may only be done by a certified service technician. You should only perform troubleshooting and simple repairs as authorized in your product documentation, or as directed by the online or telephone service and support team. Damage due to servicing that is not authorized by Dell is not covered by your warranty. Read and follow the safety instructions that are shipped with your product.

() NOTE: It is recommended that you always use a static mat and static strap while working on components in the interior of the system.

Steps

- 1. Turn off the system, including any attached peripherals, and disconnect the system from the electrical outlet and peripherals.
- 2. Open the system.
- 3. Locate the PERC card on the system board.

CAUTION: To prevent damage to the card, you must hold the card by its edges only.

- 4. Lift the card to remove it from the connector on the system board.
- 5. Disconnect the SAS cables connected to the card:
 - a. Press down and hold the metal tab on the SAS cable connector.
 - **b.** Pull the SAS cable out of the connector.
- 6. Replace the storage controller card and connect the cable. For more information on installing the card, see Install the H750 adapter SAS.
- 7. Close the system.
- 8. Reconnect the system to its electrical outlet and turn the system on, including any attached peripherals.



Figure 18. Remove PERC H750 adapter SAS

Install the PERC H750 adapter SAS

Prerequisites

CAUTION: Many repairs may only be done by a certified service technician. You should only perform troubleshooting and simple repairs as authorized in your product documentation, or as directed by the online or telephone service and support team. Damage due to servicing that is not authorized by Dell is not covered by your warranty. Read and follow the safety instructions that are shipped with your product.

Install and remove a PERC 11 card 45



NOTE: It is recommended that you always use a static mat and static strap while working on components in the interior of the system.

Steps

- 1. Turn off the system, including any attached peripherals, and disconnect the system from the electrical outlet.
- 2. Open the system.
- 3. Align the card-edge connector with the connector on the system board.

\triangle CAUTION: To prevent damage to the card, you must hold the card by its edges only.

- 4. Press the card-edge down until the card is fully seated.
- 5. Connect the SAS data cable connector to the card.

NOTE: Ensure that you connect the cable according to the connector labels on the cable. The cable does not function properly if reversed.

- 6. Route the SAS data cable through the channel on the inner side of the chassis to the backplane.
- 7. Attach the connector labeled SAS A to connector SAS A on the backplane.
- 8. Close the system.
- 9. Reconnect the system to its electrical outlet and turn the system on, including any attached peripherals.



Figure 19. Install PERC H750 adapter SAS

Remove the PERC H355 adapter SAS

Describes the tasks to remove a PERC H355 adapter SAS controller from a server.

Prerequisites

CAUTION: Many repairs may only be done by a certified service technician. You should only perform troubleshooting and simple repairs as authorized in your product documentation, or as directed by the online or telephone service and support team. Damage due to servicing that is not authorized by Dell is not covered by your warranty. Read and follow the safety instructions that are shipped with your product.

NOTE: It is recommended that you always use a static mat and static strap while working on components in the interior of the system.

Steps

1. Turn off the system, including any attached peripherals, and disconnect the system from the electrical outlet and peripherals.



- 2. Open the system.
- 3. Locate the PERC card in the expansion riser on the system board.

CAUTION: To prevent damage to the card, you must hold the card by its edges only.

- 4. Unfasten and lift the riser from the system board. Remove the PERC card.
- **5.** Disconnect any SAS cables connected to the card:
 - a. Press down and hold the metal tab on the SAS cable connector.
 - **b.** Pull the SAS cable out of the connector.
- 6. Replace the storage controller and reconnect the SAS cable before placing them in the riser. For more information on installing the card, see Install the PERC H355 adapter.
- 7. Reinstall the riser on the system board and fasten the riser.
- 8. Close the system.
- 9. Reconnect the system to its electrical outlet and turn the system on, including any attached peripherals.



Figure 20. Remove the PERC H355 adapter SAS

Install the PERC H355 adapter SAS

Describes the tasks to install a PERC H355 adapter SAS controller in a server.

Prerequisites

- CAUTION: Many repairs may only be done by a certified service technician. You should only perform troubleshooting and simple repairs as authorized in your product documentation, or as directed by the online or telephone service and support team. Damage due to servicing that is not authorized by Dell is not covered by your warranty. Read and follow the safety instructions that are shipped with your product.
- **NOTE:** It is recommended that you always use a static mat and static strap while working on components in the interior of the system.

Steps

- 1. Turn off the system, including any attached peripherals, and disconnect the system from the electrical outlet.
- 2. Open the system.

Install and remove a PERC 11 card 47



3. Align the card-edge connector with the connector on the system board.

CAUTION: To prevent damage to the card, you must hold the card by its edges only.

- 4. Press the card-edge down until the card is fully seated.
- 5. Connect the SAS data cable connectors to the card.
 - **NOTE:** Ensure that you connect the cable according to the connector labels on the cable. The cable does not function properly if reversed.
- 6. Route the SAS data cable through the channel on the inner side of the chassis to the backplane.
- 7. Attach the connector labeled SAS A to connector SAS A on the backplane, and attach the connector labeled SAS B to connector SAS B on the backplane.
- 8. Close the system.
- 9. Reconnect the system to its electrical outlet and turn the system on, including any attached peripherals.



Figure 21. Install the PERC H355 adapter SAS

Remove the PERC H355 front SAS

Prerequisites

- CAUTION: Many repairs may only be done by a certified service technician. You should only perform troubleshooting and simple repairs as authorized in your product documentation, or as directed by the online or telephone service and support team. Damage due to servicing that is not authorized by Dell is not covered by your warranty. Read and follow the safety instructions that are shipped with your product.
- **NOTE:** It is recommended that you always use a static mat and static strap while working on components in the interior of the system.

Steps

1. Turn off the system, including any attached peripherals, and disconnect the system from the electrical outlet and peripherals.

NOTE: Perform a graceful shutdown of the system to ensure data in the cache is moved to the disk before the controller is removed.



- 2. Open the system.
- 3. Locate the PERC card in the controller carrier at the front of the system.

CAUTION: To prevent damage to the card, you must hold the card by its edges only.

4. Unscrew the fasteners on the controller carrier and slide the carrier away from the backplane, disconnecting the controller from the backplane.

If you are removing a PERC H355 front SAS controller in the upside down orientation, you must remove both the backplane and the controller at the same time because of the limited clearance available:

- **a.** Uninstall all drives from the backplane.
- b. Disconnect all cables between the PERC and the backplane.
- c. Lift the backplane and PERC from the system.
- **5.** Disconnect any cables connected to the card:
 - a. Press down and hold the metal tab on the cable connector.
 - b. Pull the cables out of the connector.
- 6. Remove the PERC controller from the controller carrier.
- 7. Insert the replacement controller into the carrier and secure it with the appropriate screws.
- 8. Take the replacement storage controller and reconnect the cables before reconnecting it to the backplane.
- If you are removing a PERC H355 front SAS controller in the upside down orientation, reattach the PERC controller to the backplane first before reinstalling the backplane into the system. For more information on installing the card, see Install the PERC H355 front.
- 9. Close the system.
- **10.** Reconnect the system to its electrical outlet and turn the system on, including any attached peripherals.



Figure 22. Remove the PERC H355 front SAS



Install the PERC H355 front SAS card

Prerequisites

CAUTION: Many repairs may only be done by a certified service technician. You should only perform troubleshooting and simple repairs as authorized in your product documentation, or as directed by the online or telephone service and support team. Damage due to servicing that is not authorized by Dell is not covered by your warranty. Read and follow the safety instructions that are shipped with your product.

NOTE: It is recommended that you always use a static mat and static strap while working on components in the interior of the system.

Steps

- 1. Turn off the system, including any attached peripherals, and disconnect the system from the electrical outlet.
 - **NOTE:** Perform a graceful shutdown of the sled to ensure that data in the cache is moved to the disk before the controller is removed.
- 2. Open the system.
- 3. Connect the PERC card to the carrier and ensure that the screws are properly fastened in place.

CAUTION: To prevent damage to the card, hold the card by its edges only.

- 4. Align the carrier with the guide pins until the controller is securely seated.
- 5. Slide the card into the connector until it is fully seated in the connector. Tighten the screws on the carrier that connect to the chassis to secure the carrier.
- 6. Connect the cable connectors to the card.

NOTE: Ensure that you connect the cable according to the connector labels on the cable. The cable does not function properly if reversed.

- 7. Close the system.
- 8. Reconnect the system to its electrical outlet and turn on the system and any attached peripherals.





Figure 23. Install the PERC H755 front SAS card

Remove the PERC H350 adapter SAS

Prerequisites

CAUTION: Many repairs may only be done by a certified service technician. You should only perform troubleshooting and simple repairs as authorized in your product documentation, or as directed by the online or telephone service and support team. Damage due to servicing that is not authorized by Dell is not covered by your warranty. Read and follow the safety instructions that are shipped with your product.

NOTE: It is recommended that you always use a static mat and static strap while working on components in the interior of the system.

Steps

- 1. Turn off the system, including any attached peripherals, and disconnect the system from the electrical outlet and peripherals.
- 2. Open the system.
- **3.** Locate the PERC card on the system board.

CAUTION: To prevent damage to the card, you must hold the card by its edges only.

- 4. Lift the card to remove it from the connector on the system board.
- **5.** Disconnect the SAS cables connected to the card:
 - **a.** Press down and hold the metal tab on the SAS cable connector.
 - **b.** Pull the SAS cable out of the connector.

Install and remove a PERC 11 card 51



- 6. Replace the storage controller card and connect the cable. For more information on installing the card, see Install the PERC H350 adapter.
- 7. Close the system.
- 8. Reconnect the system to its electrical outlet and turn the system on, including any attached peripherals.



Figure 24. Remove the PERC H350 adapter SAS

Install the PERC H350 adapter SAS

Prerequisites

- CAUTION: Many repairs may only be done by a certified service technician. You should only perform troubleshooting and simple repairs as authorized in your product documentation, or as directed by the online or telephone service and support team. Damage due to servicing that is not authorized by Dell is not covered by your warranty. Read and follow the safety instructions that are shipped with your product.
- **NOTE:** It is recommended that you always use a static mat and static strap while working on components in the interior of the system.

Steps

- 1. Turn off the system, including any attached peripherals, and disconnect the system from the electrical outlet.
- 2. Open the system.
- **3.** Align the card-edge connector with the connector on the system board.

\triangle CAUTION: To prevent damage to the card, you must hold the card by its edges only.

- 4. Press the card-edge down until the card is fully seated.
- 5. Connect the SAS data cable connector to the card.

NOTE: Ensure that you connect the cable according to the connector labels on the cable. The cable does not function properly if reversed.

- 6. Route the SAS data cable through the channel on the inner side of the chassis to the backplane.
- 7. Attach the connector labeled SAS A to connector SAS A on the backplane.
- 8. Close the system.
- 9. Reconnect the system to its electrical outlet and turn the system on, including any attached peripherals.




Figure 25. Install the PERC H350 adapter SAS

Remove PERC H350 Mini Monolithic SAS

Prerequisites

CAUTION: Many repairs may only be done by a certified service technician. You should only perform troubleshooting and simple repairs as authorized in your product documentation, or as directed by the online or telephone service and support team. Damage due to servicing that is not authorized by Dell is not covered by your warranty. Read and follow the safety instructions that are shipped with your product.

- 1. Using Phillips #2 screwdriver, loosen the screws that secure the storage controller cable to the connector on the system board.
- 2. Lift the storage controller cable to disconnect it from the connector on the system board.





Figure 26. Remove the cable

- 3. Lift one end of the card and angle it to disengage the card from the card holder on the system board.
- **4.** Lift the card out of the system.



Figure 27. Remove the PERC H350 Mini Monolithic SAS



Install PERC H350 Mini Monolithic SAS

Prerequisites

CAUTION: Many repairs may only be done by a certified service technician. You should only perform troubleshooting and simple repairs as authorized in your product documentation, or as directed by the online or telephone service and support team. Damage due to servicing that is not authorized by Dell is not covered by your warranty. Read and follow the safety instructions that are shipped with your product.

Steps

- 1. Angle the integrated storage controller card and align the end of the card with the storage controller card connector on the system board.
- 2. Lower the connector side of the storage controller card into the storage controller card connector on the system board.

(i) NOTE: Ensure that the slots on the system board align with the screw holes on the storage controller card connector.



Figure 28. Install PERC H350 Mini Monolithic SAS

- 3. Route the storage controller card cable along with the wall of the system.
- 4. Align the screws on the integrated storage controller card cable with the screw holes on the connector.
- **5.** Using Phillips #2 screwdriver, tighten the screws to secure the integrated storage controller card cable to the card connector on the system board.





Figure 29. Install the cable

56 Install and remove a PERC 11 card



Driver support for PERC 11

The PERC 11 cards require software drivers to operate with the supported operating systems.

This chapter contains the procedures for installing the drivers for the PERC 11 cards.

NOTE: The driver for PERC 11 for VMware ESXi is packaged within the VMware ESXi ISO image that is downloaded from Dell. For more information, see the VMware documentation at Virtualization Solutions Documentation. It is not recommended to have drivers from controllers prior to PERC 11 on the same system.

The two methods for installing a driver that is discussed in this chapter are:

- Installing a driver during operating system installation: Use this method if you are performing a new installation of the operating system and want to include the drivers.
- **Updating existing drivers:** Use this method if the operating system and the HBA controllers are already installed and you want to update to the latest drivers.

Topics:

- Creating the device driver media
- Windows driver installation
- Linux driver installation
- Load the driver while installing an operating system

Creating the device driver media

Use one of the following two methods to create the device driver media:

- Downloading Drivers From The Dell Support Website
- Downloading Drivers From The Dell Systems Service And Diagnostic Tools Media

Download and save PERC 11 drivers from the support site

About this task

To download drivers from the Dell Support website:

Steps

- 1. Go to the Support Site.
- 2. Enter the Service Tag of your system in the Choose by Service Tag to get started field or select Choose from a list of all Dell products.
- **3.** Select the **System Type**, **Operating System**, and **Category** from the drop-down list. The drivers that are applicable to your selection are displayed.
- 4. Download the drivers that you require to a USB drive, CD, or DVD.
- **5.** During the operating system installation, use the media that you created to load the driver. For more information about reinstalling the operating system, see the relevant section for your operating system later in this guide.

Download and save PERC 11 drivers from the Dell Systems Service and Diagnostic Tools

About this task

To download drivers from the Dell Systems Service and Diagnostic Tools media:



Steps

- Insert the Dell Systems Service and Diagnostics Tools media in your system. The Welcome to Dell Service and Diagnostic Utilities screen is displayed.
- 2. Select your system model and operating system.
- 3. Click Continue.
- 4. From the list of drivers displayed, select the driver you require.
- 5. Select the self-extracting ZIP file and click $\ensuremath{\textbf{Run}}$.
- 6. Copy the driver to a CD, DVD, or USB drive.
- 7. Repeat steps 1 to 6 for all the drivers you require.

Windows driver installation

Before you install the Windows driver for PERC 11, you must first create a device driver media.

- Read the Microsoft **Getting Started** document that shipped with your operating system.
- Ensure that your system has the latest BIOS, firmware, and driver updates. If required, download the latest BIOS, firmware, and driver updates from Support Site.
- Create a device driver media using one of the methods listed below:
 - USB drive
 - CD
 - DVD

Install PERC 11 driver while newly installing the Windows Server 2016 and later

About this task

To install the driver:

Steps

- 1. Boot the system using the Windows Server 2016, or newer media.
- 2. Follow the on-screen instructions until you reach Where do you want to install Windows Server 2016 or later window and then select Load driver.
- 3. As prompted, insert the installation media and browse to the appropriate location.
- 4. Select a PERC 11 series card from the list.
- 5. Click **Next** and continue installation.

Install PERC 11 driver on which the Windows Server 2016 is already installed and later

About this task

Perform the following steps to configure the driver for the RAID controller on which the Windows Server 2016 is already installed:

- 1. Turn off the system.
- 2. Install the new RAID controller in the system.
- For detailed instructions on installing the RAID controller in the system, see Install and remove a PERC 11 card.
- **3.** Turn on the system.
 - The Found New Hardware Wizard screen displays the detected hardware device.
- 4. Click Next



- 5. On the Locate device driver screen, select Search for a suitable driver for my device and click Next.
- ${\bf 6.}~$ Browse and select the drivers from the ${\bf Locate}~{\bf Driver}~{\bf Files}$ screen.
- 7. Click Next.
 - The wizard detects and installs the appropriate device drivers for the new RAID controller.
- 8. Click Finish to complete the installation.
- 9. Reboot the system when prompted.

Update PERC 11 driver that runs on Windows Server 2016 or later

Prerequisites

(i) NOTE: Close all applications on your system before you update the driver.

Steps

- 1. Insert the media containing the driver.
- 2. Select Start > Settings > Control Panel > System. The System Properties screen is displayed.

(i) NOTE: The path to System might vary depending on the operating system family.

- **3.** Click the **Hardware** tab.
- 4. Click Device Manager.

The **Device Manager** screen is displayed.

(i) NOTE: The path to Device Manager might vary depending on the operating system family.

- 5. Expand Storage Controllers by double-clicking the entry or by clicking on the plus (+) symbol next to Storage Controllers.
- 6. Double-click the controller for which you want to update the driver.
- 7. Click the Driver tab and click Update Driver.
- The screen to update the device driver wizard is displayed.
- 8. Select Install from a list or specific location.
- 9. Click Next.
- 10. Follow the steps in the wizard and browse to the location of the driver files.
- 11. Select the INF file from the drive media.
- 12. Click Next and continue the installation steps in the wizard.
- 13. Click Finish to exit the wizard and reboot the system for the changes to take place.
 - (i) NOTE: Dell provides the Dell Update Package (DUP) to update drivers on systems running Windows Server 2016 and newer operating system. DUP is an executable application that updates drivers for specific devices. DUP supports command line interface and silent execution. For more information, see Support Site.

Linux driver installation

The driver update disk (DUD) images are created only for those operating system releases in which the native (in-box) driver is insufficient for installation. In the event that an operating system is being installed with a corresponding DUD image, see, Installing or updating the RPM driver package with KMOD support. If not, proceed with using the native device driver and then skip to the topic Installing or Updating the RPM Driver Package With KMP Support.

(i) NOTE: The driver update disk (DUD) images are created only for those operating system releases in which the native (in-box) driver is insufficient for installation. In the event that an operating system is being installed with a corresponding DUD image, follow the instructions below.

(i) NOTE: To view the complete list of boot loader options, see the installation guide of your operating system.

NOTE: If using out-of-box drivers with RHEL 7 and higher, a tainted kernel message will be displayed in the log. RedHat does not provide a mechanism to sign external drivers for RHEL.



Install or update a RPM driver package using the KMOD support

Prerequisites

(i) NOTE: This procedure is applicable for Red Hat Enterprise Linux 7.x and higher.

About this task

Perform the following steps to install the RPM package with KMOD support:

Steps

- 1. Uncompress the gzipped tarball driver release package.
- 2. Install the driver package using the command: rpm -ihv kmodmegaraid_ sas-<version>.rpm.

(i) NOTE: Use rpm -Uvh <package name> when upgrading an existing package.

- 3. If the previous device driver is in use, you must reboot the system for the updated driver to take effect.
- 4. Verify the loaded driver version by running the following command: modinfo megaraid_sas.

Install or update a RPM driver package using the KMP support

Prerequisites

(i) NOTE: This procedure is applicable for SUSE Enterprise Linux 15.x.

About this task

Perform the following steps to install the RPM package with KMP support:

Steps

- 1. Uncompress the gzipped tarball driver release package.
- 2. Install the driver package using the command: rpm -ihv kmpmegaraid_ sas- <version>.rpm.

(i) NOTE: Use rpm -Uvh <package name> when updating an existing package.

- 3. If the previous device driver is in use, you must reboot the system for the updated driver to take effect.
- 4. Verify the loaded driver version by running the following command: modinfo megaraid_sas.

Upgrading the Kernel

About this task

When upgrading to a new kernel, you must reinstall the DKMS-enabled driver packages. Perform the following steps to update or install the driver for a new kernel:

- 1. At a terminal window, type the following: dkms build -m <module_name> v <module version> k <kernel version> dkms install -m <module name> - v <module version> - k <kernel version>.
- 2. To check if the driver is successfully installed in the new kernel, type: dkms status. A message similar to the following is displayed: <driver name>, <driver version>, <new kernel version>: installed.
- 3. If the previous device driver is in use, you must restart the system for the updated driver to take effect.



Load the driver while installing an operating system

Steps

- 1. Perform the following operations to install driver media:
 - PERC Linux driver ISO:
 - a. Download the PERC Linux driver package from the Dell Support site.
 - b. Extract two base directories from the tar.gz package (tar.gz > tar > base directories).
 - c. Extract the ISO file that is available in the zipped disks-x directory. For example, RHEL79/disks-1/ megaraid_sas-07.719.03.00_el7.9-1.x86_64.iso.gz > megaraid_sas-07.719.03.00_el7.9-1.x86_64.iso
 - **d.** Mount the ISO to the Server, burn the ISO to a CD or DVD or copy the ISO file to a USB. The USB has to match with the ISO.
 - LC driver pack:
 - a. Install the LC driver pack.
 - **b.** Boot the life-cycle controller and go through the operating system deployment wizard.
- 2. Boot to the installer.
- 3. In the Installation screen, press E.
- **4.** Perform the following operation:
 - If the operating system is Red Hat Enterprise Linux 7 or RHEL 8, the CLI displays the syntax vmlinuz. Enter inst.dd.

For example, when you are prompted with the command vmlinuz intrd=initrd.img inst.stage2=hd:LABEL=RHEL-7.0\x20x86_64 quiet inst.dd.

• If the operating system is SLES 15, the CLI displays the syntax linuxefi.. Enter **dud=1**.

For example, when you are prompted with the command linuxefi/boot/x86_64/loader/linux splash=silent dud=1.

NOTE: Boot parameters may vary based on the operating system version. See operating system installation manuals for exact boot parameter syntax.

- 5. Attach the driver media (ISO, USB).
- **6.** Press F10 to boot to the operating system. A screen is displayed prompting you to select the driver media (USB, CD, ISO, and so on).
- 7. When prompted, select the driver media.

If applicable select the PERC driver ...megaraid_sas...

(i) NOTE: Ensure that the driver is selected with an X symbol.

- 8. The driver should be extracted or loaded.
- 9. Before proceeding or exiting the driver select menu, disconnect the driver media.

NOTE: Ensure that you disconnect the driver media so that the drivers are loaded successfully. If the installation media is deleted, reattach it.

10. Press C or exit to go to the installation.





This section provides information about downloading and installing the firmware using Dell Update Package (DUP).

Topics:

• Upgrade firmware controller using Dell Update Package (DUP)

Upgrade firmware controller using Dell Update Package (DUP)

About this task

(i) NOTE: If the Online Capacity Expansion operation is in progress then you cannot update the firmware version.

Steps

- 1. Go to the Drivers and Downloads page on the support site.
- 2. Locate your controller.
- 3. Download the DUP file.
 - a. To upgrade by using Windows or iDRAC, download the Windows executable file.
 - **b.** To upgrade using Linux, download the **.bin** file.

(i) NOTE: For VMware, firmware must be upgraded by using iDRAC or the PERC CLI.

- 4. Install the DUP by doing one of the following:
 - a. For Windows, run the executable file in the Windows environment.
 - **b.** For Linux, run the **.bin** file in the Linux environment.
 - c. For iDRAC, click System iDRAC > Maintenance > System Update, upload Windows executable, and then install.



Manage PERC 11 controllers using HII configuration utility

The Human Interface Infrastructure (HII) configuration utility is a storage management application integrated into the System BIOS < F2 >. It is used to configure and manage the controller(s), virtual disks, and physical disks. This utility is independent of the operating system.

Topics:

- Enter the PERC 11 HII configuration utility
- Exit the PERC 11 HII configuration utility
- Navigate to Dell PERC 11 configuration utility
- View the HII Configuration utility dashboard
- Configuration management
- Controller management
- Virtual disk management
- Physical disk management
- Hardware components
- Security key management in HII configuration utility

Enter the PERC 11 HII configuration utility

About this task

Perform the following steps to boot to the HII configuration utility:

Steps

- 1. Turn on the system.
- 2. While the system startup, press <F2> to enter System Setup.

3. Click Device Settings.

Device Settings screen lists all the RAID controllers in the system.

To access the management menu for the controller, use the arrow keys or the mouse.

- () NOTE: For more information in all the options, click Help that is available on the top right-hand corner of the browser screen. Help information for individual option menus can also be viewed by scrolling down on each option.
- **NOTE:** Some of the options within the HII configuration utility are not present if the controller does not support the corresponding feature. Options may also be grayed out if the feature is not applicable to the current configuration.

Exit the PERC 11 HII configuration utility

About this task

To exit the HII configuration utility, perform the following steps:

- 1. Click **Finish** at the bottom-right corner on the **System Setup Main Menu** screen. Displays a warning message to confirm your choice.
- 2. Click Yes to exit the HII configuration utility.



Navigate to Dell PERC 11 configuration utility

Steps

- 1. Enter the UEFI configuration Utility. See Enter the PERC 11 HII configuration utility. The **Device Settings** screen displays a list of NIC ports and the RAID controllers.
- To enter PERC 11 configuration utility, click the appropriate PERC controllers. The Dashboard view screen is displayed.

View the HII Configuration utility dashboard

The first screen that is displayed when you access the HII Configuration Utility is the **Dashboard View** screen. The following table provides detailed information about the options available on the **Dashboard View** screen.

Table 9. Dashboard view screen

Dashboard view options	Description
Main menu	Displays the following configuration options: • Configuration Management • Controller Management • Virtual Disk Management • Physical Disk Management • Hardware Components
Help	Provides context sensitive help message.
Properties	 Displays the following information about the controller: Status — displays the status of the controller. Backplane — displays information about the number of backplanes connected to the controller. BBU — displays information about the availability of Battery Backup Unit (BBU). Enclosure — displays information about the number of enclosures connected to the controller. Physical Disks — displays information about the number of physical disks connected to the controller. Disk Groups — displays information about the number of disk groups connected to the controller. Virtual Disks — displays information about the number of virtual disks connected to the controller.
View server profile	 Displays HII Spec version supported on the system and also displays the following menu options for controller components: Controller Management Hardware Components Physical Disk Management Virtual Disk Management
Actions	 Displays the following options: Configure — displays configuration options that are supported by the controller. Set Factory Defaults — restore factory default values for all controller properties.
Background operations	Displays if virtual disk or physical disk operations are in progress.



Configuration management

Auto Configure RAID 0

Steps

- 1. Enter the Dell PERC 11 Configuration Utility. See Navigate to Dell PERC 11 configuration utility.
- 2. Click Main Menu > Configuration Management > Auto Configure RAID 0.
- **3.** Select **Confirm** and click **Yes** to continue. A RAID 0 Virtual disk is created on all physical disks that are in Ready state.

Create virtual disks

Steps

- 1. Enter the Dell PERC 11 Configuration Utility. See, Navigate to Dell PERC 11 configuration utility.
- 2. Click Main Menu > Configuration Management > Create Virtual Disk. The following list of options are displayed for you to define the virtual disk parameters:

Table 10. Create virtual disks

Option	Description
Create Virtual Disk	Allows you to create virtual disk selecting the RAID level, physical disks, and virtual disk parameters
Select RAID level	Allows you to choose the RAID level of your choice
Secure Virtual Disk	If you want to create a secured virtual disk, select Secure Virtual Disk. (i) NOTE: The Secure Virtual Disk option is enabled by default, only if the security key has been configured. Only SED physical disks are listed.
Select Physical Disks From	 Allows you to select one of the physical disk capacities: Unconfigured Capacity: creates a virtual disk on unconfigured physical disks. Free Capacity: utilizes unused physical disk capacity that is already part of a disk group.
Select Physical Disks	If you want to select the physical disks from which the virtual disks are being created, click Select Physical Disks . This option is displayed if you select Unconfigured Capacity as your physical disk capacity.
Select Disk Groups	If you want to select the disk groups from which the virtual disks are being created, click Select Disk Group . This option is displayed if you select Free Capacity as your physical disk capacity.
Configure Virtual Disk Parameters	Allows you to set the virtual disk parameters when creating the virtual disk. For more information, see Configuring virtual disk parameters.

3. Click Create Virtual Disk.

The virtual disk is created successfully.

NOTE: Ensure that you restart the system after creating a new Non-RAID or Virtual Disk on drives that previously had boot partitions.



Configure virtual disk parameters

Steps

- 1. Create a virtual disk, see Creating the virtual disks.
 - The Configure Virtual Disk Parameters section is displayed on the Create Virtual Disk screen.
- 2. In the Configure Virtual Disk Parameters section, you can set the following virtual disk parameters:

Table 11. Configure virtual disk parameters

Virtual disk parameters	Description
Virtual Disk Name	Allows you to enter the name for the virtual disk (i) NOTE: Allowed characters are A-Z, a-z, 0-9, underscore (_), and hyphen (-) only.
Virtual Disk Size	Displays the maximum capacity available for the virtual disk
Virtual Disk Size Unit	Displays the virtual disk storage space in megabytes, gigabytes, and terabyte.
Strip Element Size	Allows you to select the strip element size The disk striping involves partitioning each physical disk storage space in stripes of the following sizes: 64 KB, 128 KB, 256 KB, 512 KB, and 1 MB. By default, the strip element size is set to 256 KB.
Read Policy	 Displays the controller read policy You can set the read policy to: No read ahead—specifies that the controller does not use read ahead for the current virtual disk. Read ahead—specifies that the controller uses read ahead for the current virtual disk. Read ahead capability allows the controller to read sequentially ahead of requested data and store the additional data in the cache memory, anticipating that the data is required soon. By default, the read cache policy is set to read ahead.
Write Policy	 Displays the controller write cache policy You can set the write policy to: Write through—the controller sends a data transfer completion signal to the host when the disk subsystem has received all the data in a transaction. Write back—the controller sends a data transfer completion signal to the host when the controller cache has received all the data in a transaction. By default, the write policy is set to Write Back.
Disk Cache	Allows you to set the disk cache policy to default, enable, or disable. By default, the disk cache is set to default.
Default Initialization	 Displays the virtual disk initialization options. You can set the default initialization to: No — The virtual disk is not initialized. Fast — The first 8 MB of the virtual disk is initialized. Full — The entire virtual disk is initialized. For more information, see Virtual disk initialization. By default, the default initialization is set to No.

Create profile based virtual disk

- 1. Enter the Dell PERC 11 Configuration Utility. See Navigate to Dell PERC 11 configuration utility.
- 2. Click Main Menu > Configuration Management > Creating Profile Based Virtual Disk. The following list of RAID modes are displayed:
 - Generic RAID 0
 - Generic RAID 1
 - Generic RAID 5
 - Generic RAID 6
 - File Server



- Web/Generic Server
- Database
- **3.** Based on the RAID mode selected, one or more the physical disk selection criteria is displayed.
- **4.** From the **Physical Disk Selection Criteria** drop-down box, select a criterion based your requirement. The Profile Parameters of the selected option is displayed.
- 5. Click Create Virtual Disk
- Select Confirm and click Yes to continue. The virtual disk is created with the parameters of the profile selected.

View disk group properties

Steps

- 1. Enter the **Dell PERC 11 Configuration Utility**. See Navigate to Dell PERC 11 configuration utility.
- Click Main Menu > Configuration Management > View Disk Group Properties. The list of disk group properties are displayed:

Table 12. View disk group properties

Properties	Descriptions
Capacity Allocation	Displays all the virtual disks associated with the specific disk group. It also provides information about the available free space
Secured	Displays whether the disk group is secured or not

Convert to Non-RAID disk

Prerequisites

To convert a physical disk to non–RAID disk from the HII Configuration Utility, perform the following steps:

Steps

- 1. Enter the Dell PERC 11 Configuration Utility. See Navigate to Dell PERC 11 configuration utility.
- Click Main Menu > Configuration Management > Convert to Non-RAID Disk. The list of physical disks appears.
- 3. Select the physical disk to convert to Non-RAID disk.
- 4. Click Ok.
- A screen appears asking if you are sure you want to perform the operation.
- 5. Select the Confirm option.
- 6. Click Yes.

The operation is successful.

Delete configurations

Steps

- 1. Enter the Dell PERC 11 Configuration Utility. See Navigate to Dell PERC 11 configuration utility.
- 2. Click Main Menu > Configuration Management > Clear Configuration. A screen is displayed asking if you are sure you want to perform the operation.
- 3. CAUTION: It is recommended that you back up data stored on the virtual disks and hot spare disks on the controller before deleting the virtual drive.

Select **Confirm** and click **Yes** to continue. The virtual disks and hot spare disks available on the controller are deleted successfully.



Controller management

Clear controller events

Steps

- 1. Enter the Dell PERC 11 Configuration Utility. See Navigate to Dell PERC 11 configuration utility.
- 2. Click Main Menu > Controller Management > Advanced Controller Management.
- Click Clear Controller Events.
 A screen is displayed asking if you are sure you want to clear the controller events.
- 4. Select Confirm and click Yes to continue.

Save controller events

Steps

- 1. Enter the Dell PERC 11 Configuration Utility. See Navigate to Dell PERC 11 configuration utility.
- 2. Click Main Menu > Controller Management > Advanced Controller Management.
- Click Save Controller Events.
 A screen is displayed asking if you want to replace the existing file name.
- 4. Select Confirm and click Yes to continue.

Save debug log

Steps

- 1. Enter the Dell PERC 11 Configuration Utility. See Navigate to Dell PERC 11 configuration utility.
- 2. Click Main Menu > Controller Management > Advanced Controller Management.
- 3. Click Save Debug Log.
 - A screen is displayed indicating that the operation is successful.
- 4. Click **Ok**.

Enable security

Steps

- 1. Enter the Dell PERC 11 Configuration Utility. See Navigate to Dell PERC 11 configuration utility.
- 2. Click Main Menu > Controller Management > Advanced Controller Management.
- 3. Click Enable security, select Local Key Management.
- 4. Click Ok.
- 5. If you want to use the passphrase generated by the controller, click **Suggest Passphrase** and **Confirm** the passphrase by re-entering.
 - The operation is successful.
- Select I Recorded the Security Settings For Future Reference, click Enable Security. A screen is displayed indicating that the security will be enabled on this controller if you proceed.
- 7. Select **Confirm** and click **Yes** to continue. The operation is successful and click **Ok**.

Disable security

Steps

1. Enter the Dell PERC 11 Configuration Utility. See Navigate to Dell PERC 11 configuration utility.



- 2. Click Main Menu > Controller Management > Advanced Controller Management.
- 3. Click Disable security.
 - A screen is displayed asking if you are sure you want to disable security.
- Select Confirm and click Yes to continue. The operation is successful and click Ok.

Change security settings

Steps

- 1. Enter the Dell PERC 11 Configuration Utility. See Navigate to Dell PERC 11 configuration utility.
- 2. Click Main Menu > Controller Management > Advanced Controller Management.
- 3. Click Change Security Settings, select Change Current Security Settings.
- 4. Click Ok.
- 5. If you want to use the passphrase generated by the controller, click **Suggest Passphrase** and **Confirm** the passphrase by re-entering.
- The operation is successful.
- 6. Click Save Security Settings.
- Select Confirm and click Yes to continue. The operation is successful and click Ok.

Restore factory default settings

Steps

- 1. Enter the Dell PERC 11 Configuration Utility. See Navigate to Dell PERC 11 configuration utility.
- 2. Click Main Menu > Set Factory Defaults.
- A screen is displayed asking you to confirm the operation.
- 3. Select Confirm and click Yes to continue.

Auto configure behavior

Steps

- 1. Enter the Dell PERC 11 Configuration Utility. See Navigate to Dell PERC 11 configuration utility.
- 2. Click Main Menu > Controller Management > Advanced Controller Management > Manage Controller Mode. You can view the current Controller Mode.
- 3. Click Manage Controller Mode.
 - If required, you can view or change the hard drive settings for the controller. The possible options are:
 - Off and Non-RAID Disk
- 4. Click Apply Changes to save the changes.
- 5. Select Confirm and click Yes to continue.

NOTE: This feature is supported on PERC H355 adapter SAS, PERC H355 front SAS, PERC H350 Mini Monolithic SAS, and PERC H350 adapter SAS.

Manage controller profile

About this task

View the details of the profile and choose the desired profile, if supported. To view the properties of the controller profiles:

Steps

1. Enter the Dell PERC 11 Configuration Utility. See Navigate to Dell PERC 11 configuration utility.

Manage PERC 11 controllers using HII configuration utility 69



2. Click Main Menu > Controller Management > Advanced Controller Management > Manage Controller Profiles. The current profile and profile properties are displayed.

Advanced controller properties

Set the patrol read mode

Steps

- 1. Enter the Dell PERC 11 Configuration Utility. See Navigate to Dell PERC 11 configuration utility.
- 2. Click Main Menu > Controller Management > Advanced Controller Properties.
- 3. Click Patrol Read.
 - The following options are displayed:
 - Start—Starts patrol read for the selected controller.
 - Suspend—Suspends the ongoing patrol read operation on the controller.
 - Resume—Resumes the suspended patrol read operation.
 - Stop—Stops patrol read for the selected controller.
- 4. Set the Mode to Auto, Manual, or Disabled.
- 5. Click Apply Changes.

Enable physical disk power management

Steps

- 1. Enter the Dell PERC 11 Configuration Utility. See Navigate to Dell PERC 11 configuration utility.
- 2. Click Main Menu > Controller Management > Advanced Controller Properties.
- 3. Click Physical Disk Power Management.

The following list of options is displayed:

- Time Interval for Spin Down—allows the user to specify the delay time before a disk is spun down.
- Spin Down Hot Spare—allows you to enable or disable the spin down of hot spare disks.
- Spin Down Unconfigured Good—spin down of un-configured disks.
- Select the applicable options and click Apply Changes. The changes made are saved successfully.

Configure hot spare drives

Steps

- 1. Enter the Dell PERC 11 Configuration Utility. See Navigate to Dell PERC 11 configuration utility.
- 2. Click Main Menu > Controller Management > Advanced Controller Properties.
- 3. Click Spare.
 - The following list of options are displayed:
 - Persistent Hot Spare—allows you to enable or disable the ability to have same system backplane or storage enclosure disk slots dedicated as hot spare slots.
 - Allow Replace Member with Revertible Hot Spare—allows you to enable or disable the option to copy the data form a hot spare disk to physical disk.
 - Auto Replace Member on Predictive Failure—allows you to enable or disable the option to start a Replace Member operation if a predictive failure error is detected on a physical disk.
- 4. Select the applicable option and click **Apply Changes**.

The changes made are saved successfully.



Set task rates

Steps

- 1. Enter the Dell PERC 11 Configuration Utility. See Navigate to Dell PERC 11 configuration utility.
- 2. Click Main Menu > Controller Management > Advanced Controller Properties.
- 3. Click Task Rates.
 - The following options are displayed:
 - Background Initialization (BGI) Rate
 - Consistency Check Rate
 - Rebuild Rate
 - Reconstruction Rate
- 4. You can make the necessary changes and then click **Apply Changes**. The task rates operation is completely successfully.

Properties of Enterprise Key Management (EKM)

Steps

- 1. Enter the Dell PERC 11 Configuration Utility. See Navigate to Dell PERC 11 configuration utility.
- 2. Click Main Menu > Controller Management > Advanced Controller Properties.
- **3.** Click **Enterprise Key Management**. The properties of Enterprise Key Management is displayed.

Controller properties

Auto import foreign configuration

Steps

- 1. Enter the Dell PERC 11 Configuration Utility. See Navigate to Dell PERC 11 configuration utility.
- 2. Click Main Menu > Controller Management > Advanced Controller Properties.
- 3. In the Controller Properties section, set the Auto Import Foreign Configuration option to Enabled or Disabled.
- 4. Click Apply Changes.

Disable auto import

Steps

- 1. Enter the Dell PERC 11 Configuration Utility. See Navigate to Dell PERC 11 configuration utility.
- 2. Click Main Menu > Controller Management > Advanced Controller Properties.
- 3. In the Controller Properties section, set the Auto Import Foreign Configuration option to Disabled.
- **4.** Click **Apply Changes**. The auto import is disabled successfully.

Enable auto import

- 1. Enter the Dell PERC 11 Configuration Utility. See Navigate to Dell PERC 11 configuration utility.
- 2. Click Main Menu > Controller Management > Advanced Controller Properties.
- 3. In the Controller Properties section, set the Auto Import Foreign Configuration option to Enabled.
- 4. Click Apply Changes. The auto import is enabled successfully.



Select boot mode

Steps

- 1. Enter the Dell PERC 11 Configuration Utility. See Navigate to Dell PERC 11 configuration utility.
- 2. Click Main Menu > Controller Management > Advanced Controller Properties.
- **3.** In the **Controller Properties** section, select boot mode from the **Boot Mode** drop-down box. The following lists of boot mode options appear:

Table 13. Boot mode options

Option	Description
Stop on errors	The system stops during boot for errors which require attention from the user to rectify the issue.
Pause on errors	System pauses during boot to show errors but continue boot after it times out. Only critical events with an infinite timeout halt boot and require the user's attention to correct the issue.

NOTE: In UEFI BIOS mode, errors with timeouts do not appear during boot. It is designed to arise only in legacy BIOS mode.

(i) NOTE: By default, the boot mode option is set to pause on errors.

4. Click Apply Changes.

The boot mode operation is completed successfully.

Abort the consistency check

Steps

- 1. Enter the Dell PERC 11 Configuration Utility. See Navigate to Dell PERC 11 configuration utility.
- 2. Click Main Menu > Controller Management > Advanced Controller Properties.
- 3. In the Controller Properties section, set the Abort Consistency Check on Error option to Enabled or Disabled.
- 4. Click Apply Changes.
- The option to abort the consistency check operation on a redundant virtual disk is enabled if there is any inconsistency found in the data.

Preboot trace buffer

Steps

- 1. Enter the Dell PERC 11 Configuration Utility. See Navigate to Dell PERC 11 configuration utility.
- 2. Click Main Menu > Controller Management > Advanced Controller Properties.
- 3. In the Controller Properties section, set the Preboot Trace Buffer option to Enabled or Disabled.
- 4. Click Apply Changes.

Clear the cache memory

- 1. Enter the Dell PERC 11 Configuration Utility. See Navigate to Dell PERC 11 configuration utility.
- 2. Click Main Menu > Controller Management > Advanced Controller Properties.
- **3.** Click **Cache and Memory** > **Discard Preserved Cache**. The preserved cache is cleared successfully.
- 72 Manage PERC 11 controllers using HII configuration utility



Enable boot support

Steps

- 1. Enter the Dell PERC 11 Configuration Utility. See Navigate to Dell PERC 11 configuration utility.
- 2. Click Main Menu > Controller Management.
- 3. From the Select Boot Device drop-down box, select the primary bootable device.

In **Select Boot Device**, you will not be able to view 4 K sector drives. To view all the virtual disks created, navigate to the **Virtual Disk Management** screen in HII. For more information, see Virtual disk management.

If no boot device is selected, the first virtual disk will be set as the boot device on the next reboot. A Non-RAID disk is auto-selected as the boot device, if the controller does not have any virtual disks present.

- (i) **NOTE: Select Boot Device** is only applicable in legacy BIOS mode.
- (i) NOTE: 4 K sector drives boot support is only available in UEFI mode and managed by the boot loader.
- 4. Click Apply Changes.

Boot support is enabled for the selected controller.

Virtual disk management

Virtual disk numbering

Virtual disks are numbered in descending order beginning with the highest, ID 239.

View virtual disk properties

Steps

- 1. Enter the Dell PERC 11 Configuration Utility. See Navigate to Dell PERC 11 configuration utility.
- Click Main Menu > Virtual Disk Management.
 All the virtual disks associated with the RAID controller are displayed.
- **3.** To view the properties, click on the virtual disk. You can view the following properties of the Virtual disk:

Table 14. Virtual disk properties

Option	Description
Operation	List of operations you can perform on the selected virtual disk. The options are: Blink Unblink Delete Virtual Disk Reconfigure Virtual Disks Fast Initialization Slow Initialization
Name	Indicates the name of the virtual disk.
RAID level	Indicates the RAID level of the virtual disk.
Status	Indicates the status of the virtual disk. The possible options are: Optimal Degraded Offline Failed
Size	Indicates the size of the virtual disk.



View physical disks associated with a virtual disk

Steps

- 1. Enter the Dell PERC 11 Configuration Utility. See Navigate to Dell PERC 11 configuration utility.
- 2. Click Main Menu > Virtual Disk Management. All the virtual disks associated with the RAID controller are displayed.
- **3.** Click on a virtual disk. The properties of the virtual disk are displayed.
- Click View Associated Physical Disks. All the physical disks that are associated with the virtual disk are displayed.
- 5. From the **Associated Physical Disks** section, select the physical disk.
- 6. Click View Physical Disk Properties to view the physical disk properties.

View advanced properties of a virtual disk

Steps

- 1. Enter the Dell PERC 11 Configuration Utility. See Navigate to Dell PERC 11 configuration utility.
- 2. Click Main Menu > Virtual Disk Management.

All the virtual disks associated with the RAID controller are displayed.

- **3.** Click the virtual disk. The properties of the virtual disk are displayed.
- 4. Click Advanced....
 - You can view the following additional properties of the virtual disk:

Table 15. Advanced properties of the virtual disk

Option	Description
Logical sector size	Indicates the logical sector size of this virtual disk.
Strip element size	Indicates the strip element size for the virtual disk.
Secured	Indicates whether the virtual disk is secured or not.
Bad blocks	Indicates whether the virtual disk has corrupted blocks.

Configure virtual disk policies

Steps

- 1. Enter the Dell PERC 11 Configuration Utility. See Navigate to Dell PERC 11 configuration utility.
- 2. Click Main Menu > Virtual Disk Management.
 - All the virtual disks associated with the RAID controller are displayed.
- 3. Click Advanced....
 - You can view the following virtual disk policies:

Table 16. Virtual disk policies

Option	Description
Current write cache	Indicates the current write cache policy for the virtual disk.
Default write cache	 Allows selection of the write cache policy for the virtual disk. The possible options are: Write Through Write Back Force Write Back
Read cache policy	Allows selection of the read cache policy for the virtual disk. The possible options are:

74 Manage PERC 11 controllers using HII configuration utility



Table 16. Virtual disk policies (continued)

Option	Description
	No Read AheadRead Ahead
Disk cache	 Allows selection of the disk cache policy for the virtual disk. The possible options are: Default (Disk Default) Enable Disable

4. Click Apply Changes.

The changes made are saved successfully.

Configure Virtual Disks

When configuring the virtual disks, you should consider the workload intended; RAID 1: for simple boot disk; RAID 5 or 6: for file or web servers (sequential reads/writes of files); RAID 10: for transactional database (small random reads and writes).

Virtual disks configured on hard drives should use the controller default cache setting of Write Back and Read Ahead.

Virtual disks configured on SSDs can use the same controller defaults settings as hard drives. Most users perform a copy of OS files or a data base to the new array. This setting provides optimum performance in this configuration.

Once the copy is complete, the array can be used as it is depending on the number and type of SSDs. It is recommended to enable FastPath by changing the controller's Write cache policy to Write Through and the Read cache policy to No Read Ahead. FastPath is developed to achieve the best random read/write performance from SSDs.

Only IO block sizes smaller than the virtual disk's stripe size are eligible for FastPath. In addition, there should be no background operations (rebuild, initialization) running on the virtual disks. FastPath is disabled if there is active background operation.

(i) NOTE: RAID 50, and RAID 60 virtual disks cannot use FastPath.

(i) NOTE: The Physical Disk Power Management feature is not applicable to FastPath-capable virtual disks.

Perform expand virtual disk operation

Prerequisites

To enable expand virtual disk feature from the HII Configuration Utility, perform the following steps:

Steps

- 1. Enter the Dell PERC 11 Configuration Utility. See Navigate to Dell PERC 11 configuration utility.
- 2. Click Main Menu > Virtual Disk Management.
- The list of virtual disks is displayed.
- **3.** Select the virtual disk.
- 4. From the Operations drop-down menu, select Expand Virtual Disk.

(i) NOTE: You can view the Expand Virtual Disk feature only if there is free space available in the associated disk group.

- 5. Click Go.
- 6. To expand virtual disk, enter the percentage of available capacity, and then click **Ok**. A screen is displayed asking if you are sure you want to perform the operation.
- 7. Select the **Confirm** option.
- 8. Click Yes.

The expand virtual disk operation is completed successfully.



Perform consistency check

Prerequisites

To enable consistency check from the HII Configuration Utility, perform the following steps:

Steps

- 1. Enter the Dell PERC 11 Configuration Utility. See Navigate to Dell PERC 11 configuration utility.
- 2. Click Main Menu > Virtual Disk Management.
- The list of virtual disks is displayed.
- **3.** Select the virtual disk.

(i) NOTE: Consistency check cannot be run on RAID 0 virtual disks.

- 4. From the Operations drop-down menu, select Check Consistency.
- Click Go.
 A screen is displayed asking if you are sure you want to perform the operation.
- 6. Select the **Confirm** option.
- 7. Click Yes.
 - The consistency check operation is completed successfully.

Physical disk management

View physical disk properties

Steps

- 1. Enter the Dell PERC 11 Configuration Utility. See Navigate to Dell PERC 11 configuration utility.
- Click Main Menu > Physical Disk Management. All the physical disks that are associated with the RAID controller are displayed.
- 3. To view the properties, click the physical disk.

Table 17. Physical disk properties

Option	Description
Operation	 The list of operations you can perform on the selected physical disk. The options are: Blink Unblink Assign global hot spare Cryptographic erase Convert to non-RAID disk
Device ID	Unique identifier of the physical disk.
Backplane ID	Backplane ID in which the physical disk is located in for PERC H755 adapter, PERC H755 front SAS, PERC H755N front NVMe, PERC H750 adapter SAS, PERC H755 MX adapter, PERC H355 adapter SAS, PERC H355 front SAS, PERC H350 adapter SAS , and PERC H350 Mini Monolithic SAS
Slot number	The drive bay in which the physical disk is located for the corresponding backplane or enclosure to which the controller is connected.
Status	Status of the physical disk.
Size	Size of the physical disk.
Туре	Type of the physical disk.
Model	Model of the physical disk.

76 Manage PERC 11 controllers using HII configuration utility



Table 17. Physical disk properties (continued)

Option	Description
Serial number	Serial of the physical disk.

4. To view additional properties of the physical disk, click Advanced....

Table 18. Advanced physical disk properties

Option	Description
Logical sector size	Logical sector size of the selected physical disk
Physical sector size	Physical sector size of the selected physical disk
SMART status	SMART status of a physical disk
Revision	Firmware version of the physical disk
WWID	Unique identifier used to identify the device
Multipath	Multipath of the controller
Physical disk power state	Power condition (On or Power Save) of the physical disk
Disk cache setting	Disk cache setting (i) NOTE: Disk cache for SATA Gen3 drives is disabled by default.
Disk protocol	Type of hard disk used
Device speed	Speed of the physical disk
Negotiated link speed	Negotiated link speed of the device
PCIe capable link width	N/A for SAS/SATA drives
PCIe negotiated link width	N/A for SAS/SATA drives
Encryption capable	Encryption capability of the physical disk
Encryption supported	Encryption capability enabled at the controller level
Secured	Security status of the physical disk
Cryptographic erase capable	Cryptographic erase capability of the physical disk

Cryptographic erase

Cryptographic erase is a process to erase all data permanently on an encryption-capable and unconfigured physical disk, and reset the security attributes.

Prerequisites

- The non-RAID and virtual disks associated with the drive are deleted.
- The disks are not hot spares.

About this task

The Cryptographic erase feature is supported only on Instant Secure Erase (ISE) and Self Encrypting Drives (SED) drives.

- 1. Enter the **Dell PERC 11 Configuration Utility**. See Navigate to Dell PERC 11 configuration utility.
- 2. Click Main Menu > Physical Disk Management.
- The list of physical disks is displayed.
- **3.** Select a physical disk.
- 4. From the Operations drop-down menu, select Cryptographic Erase.



(i) NOTE: If the drive installed is ISE or SED capable only then the Cryptographic erase option is displayed.

5. Click Go.

A screen is displayed asking if you are sure you want to perform the operation.

- 6. Select the **Confirm** option.
- 7. Click Yes.

The Cryptographic erase operation is completed successfully.

Physical disk erase

Prerequisites

To use the Physical Disk Erase feature from the HII Configuration Utility, perform the following steps:

Steps

- 1. Enter the Dell PERC 11 Configuration Utility. See Navigate to Dell PERC 11 configuration utility.
- 2. Click Main Menu > Physical Disk Management. The list of physical disks is displayed.
- 3. Select a physical disk.
- 4. From the Operations drop-down menu, select Physical Disk Erase.

(i) NOTE: If the drive installed is neither SED or ISE capable, then only the Physical Disk Erase option is displayed.

5. Click Go.

A screen is displayed asking if you are sure you want to perform the operation.

- 6. Select the **Confirm** option.
- 7. Click Yes.

The physical disk erase operation is completed successfully.

Assigning a global hot spare

Prerequisites

To assign a global hot spare from the HII Configuration Utility, perform the following steps:

Steps

- 1. Enter the Dell PERC 11 Configuration Utility. See Navigate to Dell PERC 11 configuration utility.
- 2. Click Main Menu > Physical Disk Management.
- The list of physical disks is displayed.
- **3.** Select the physical disk.
- 4. From the Operations drop-down menu, select Assign Global Hot Spare.
- 5. Click Go.
 - A screen is displayed asking if you are sure you want to perform the operation.
- 6. Select the Confirm option.
- 7. Click Yes.

The global hot spare disk is created successfully.

Assigning a dedicated hot spare

Prerequisites

To assign a dedicated hot spare from the HII Configuration Utility, perform the following steps:



Steps

- 1. Enter the Dell PERC 11 Configuration Utility. See Navigate to Dell PERC 11 configuration utility.
- 2. Click Main Menu > Physical Disk Management.
- The list of physical disks is displayed.
- **3.** Select the physical disk.
- 4. From the Operations drop-down menu, select Assign Dedicated Hot Spare.
- 5. Click Go.
- A screen is displayed asking if you are sure you want to perform the operation.
- 6. Select the **Confirm** option.
- 7. Click Yes.

The dedicated hot spare disk is created successfully.

Convert to RAID capable

Prerequisites

To convert a non-RAID disk to RAID capable disk from the HII Configuration Utility, perform the following steps:

Steps

- 1. Enter the Dell PERC 11 Configuration Utility. See Navigate to Dell PERC 11 configuration utility.
- 2. Click Main Menu > Physical Disk Management. The list of physical disks appears.
- **3.** Select the physical disk.
- 4. From the Operations drop-down menu, select Convert to RAID capable.
- 5. Click Go.
- A screen appears asking if you are sure you want to perform the operation.
- $\textbf{6.} \hspace{0.1 cm} \text{Select the } \textbf{Confirm} \hspace{0.1 cm} \text{option}.$
- 7. Click Yes. The operation is successful.

Convert to Non-RAID disk

Prerequisites

To convert a physical disk to non-RAID disk from the HII Configuration Utility, perform the following steps:

- 1. Enter the Dell PERC 11 Configuration Utility. See Navigate to Dell PERC 11 configuration utility.
- 2. Click Main Menu > Physical Disk Management. The list of physical disks appears.
- 3. Select the physical disk.
- 4. From the Operations drop-down menu, select Convert to Non-Raid disk.
- 5. Click Go.
 - A screen appears asking if you are sure you want to perform the operation.
- 6. Select the Confirm option.
- 7. Click Yes. The operation is successful.



Hardware components

View battery properties

Steps

- 1. Enter the Dell PERC 11 Configuration Utility. See Navigate to Dell PERC 11 configuration utility.
- 2. Click Main Menu > Hardware Components > Battery Management.
- The battery and capacity information are displayed.
- $\ensuremath{\textbf{3}}.$ You can view the following properties of the battery:

Table 19. Battery properties

Field	Description
Туре	Displays the type of battery available.
Status	Displays the current status of the battery.
Temperature	Displays the current temperature of the battery and also indicates whether the temperature is normal or high.
Charge	Displays the available charge of the battery in percentage.

4. Displays click Advanced....

The additional advanced properties of the physical battery are displayed.

5. You can view the following advanced properties of the battery:

Table 20. Advanced battery properties

Field	Description
Status	Displays whether the current status of the battery is learning, degraded, or failed.
Voltage	Displays whether the voltage status of the battery is normal or high.
Current	Displays power consumption of the battery in milliamps (mA).
Full capacity	Displays the maximum charge capacity of the battery.
Remaining capacity	Displays the current charge capacity of the battery.
Expected margin of error	Displays expected margin of error.
Completed discharge cycles	Displays the completed discharge cycles.
Learn mode	Displays the condition of the battery. The learn cycle is a periodic operation that calculates the charge that is remaining in the battery to ensure there is sufficient energy.

View physical disks associated with an enclosure

- 1. Enter the Dell PERC 11 Configuration Utility. See Navigate to Dell PERC 11 configuration utility.
- 2. Click Main Menu > Hardware Components > Enclosure Management.
- 3. From the Select Enclosure field, choose the enclosure for which you need to view the physical disks.
- All the physical disks that are associated with the virtual disk are displayed.
- Click the Attached Physical Disks drop-down box.
 All the physical disks that are associated with the selected enclosure are displayed.
- 80 Manage PERC 11 controllers using HII configuration utility



Security key management in HII configuration utility

The Dell OpenManage storage management application and the **HII Configuration Utility** of the controller allow security keys to be created and managed as well as create secured virtual disks. The following section describes the menu options specific to security key management and provide detailed instructions to perform the configuration tasks. The contents in the following section apply to the **HII Configuration Utility**. For more information on the management applications, see Applications and User Interfaces supported by PERC 11.

- The **Controller Management** screen displays controller information and action menus. You can perform the following security-related actions through the controller management menu:
 - **Security Key Management**—Creates or changes the local key management (LKM) security key. Deletes the local key management (LKM) or secure enterprise key manager (SEKM) security key.
- The **Virtual Disk Management** screen displays physical disk information and action menus. You can perform the following security related actions through the virtual disk management menu:
 - Secure Disk Group—Secures all virtual disks in disk group.
 - **Create secure virtual disk**—Creates a new virtual disk that is secured with the security key on the controller.
- The **Physical Disk Management** screen displays physical disk information and action menus. You can perform the following security-related actions through the physical disk management menu:
 - \circ $\,$ Secure non-RAID disk—Secures the non-RAID disk with the controller security key.
 - **Cryptographic Erase**—Permanently erases all data on the physical disk and resets the security attributes.

For more information on the Physical Disk Management screen and the Virtual Disk Management screen, see Physical disk management and Virtual disk management.



Security key and RAID management

Topics:

- Security key implementation
- Local Key Management
- Create a security key
- Change Security Settings
- Disable security key
- Create a secured virtual disk
- Secure a non-RAID disk
- Secure a pre-existing virtual disk
- Import a secured non-RAID disk
- Import a secured virtual disk
- Dell Technologies OpenManage Secure Enterprise Key Manager

Security key implementation

The PERC 11 series of cards support self-encrypting disk (SED) for protection of data against loss or theft of SEDs. Protection is achieved by the use of encryption technology on the drives. There is one security key per controller. You can manage the security key using local key management (LKM) or OpenManage Secure Enterprise Key Manager, also referred as Secure Enterprise Key Manager (SEKM). The LKM key can be escrowed in to a file using Dell OpenManage Storage Management application. The security key is used by the controller to lock and unlock access to encryption-capable physical disks. In order to take advantage of this feature, you must:

- **1.** Have SEDs in your system.
- **2.** Create a security key.
- **NOTE:** If the host system is powered off when connected to an external enclosures or if the sled is powered off in C6XXX PowerEdge servers, the drives will remain in an unlocked state until they are power cycled or AC power is disconnected from the sled or external enclosure.

Local Key Management

You can use Local Key Management (LKM) to generate the key ID and the passphrase that is required to secure the virtual disk. You can secure virtual disks, change security keys, and manage secured foreign configurations using this security mode. NOTE: LKM mode is not supported on PERC H355 adapter SAS, PERC H350 adapter SAS, PERC H355 front SAS, and PERC H350 Mini Monolithic SAS.

Create a security key

About this task

(i) NOTE: There is no passphrase backup option when you create a security key; you need to remember your passphrase.

Steps

- 1. Enter the Dell PERC 11 Configuration Utility. See Navigate to Dell PERC 11 configuration utility.
- 2. Click Main Menu > Controller Management > Advanced Controller Management > Enable Security.
- 3. Select the Security Key Management mode as Local Key Management.

Inserido ao protocolo **22.951.206-4** por: **Pedro Henrique Golin Linhares** em: 31/03/2025 12:49. A autenticidade deste documento pode ser validada no endereço: https://www.eprotocolo.pr.gov.br/spiweb/validarDocumento com o código: fc1c44dfe0686d5547ffd691199be699.



4. Click Ok.

5. In the Security Key Identifier field, enter an identifier for your security key.

NOTE: The Security Key Identifier is a user supplied clear text label used to associate the correct security key with the controller.

- 6. If you want to use the passphrase generated by the controller, click **Suggest Passphrase**. Assigns a passphrase suggested by the controller automatically.
- 7. In the **Passphrase** field, enter the passphrase.

NOTE: Passphrase is case-sensitive. You must enter minimum 8 or maximum 32 characters. Ensure that the characters contain at least one number, one lower case letter, one upper case letter, and one non-alphanumeric character.

8. In the **Confirm** field, re-enter the passphrase to confirm.

NOTE: If the Passphrase entered in the Passphrase and Confirm fields do not match, then you are prompted with an error message to enter the passphrase again.

- 9. Select the I recorded the Security Settings for Future Reference option.
- 10. Click Enable Security.

The Security Key is created successfully.

Change Security Settings

Steps

- 1. Enter the Dell PERC 11 Configuration Utility. See Navigate to Dell PERC 11 configuration utility.
- 2. Click Main Menu > Controller Management > Advanced Controller Management > Change Security Settings.
- 3. Select security identifier:
 - a. To change the Security key Identifier enter a new key identifier in Enter a New Security Key identifier text box.
 - b. To keep existing key identifier, select Use the existing Security Key Identifier check box.
- 4. Enter the existing passphrase.
- 5. Set passphrase:
 - a. To change the security passphrase, enter a new passphrase in the **Enter a New Passphrase** text box. Re-enter the new passphrase to confirm.
 - b. To keep the existing passphrase, select Use the existing passphrase.
- 6. Select I recorded the Security Settings for Future Reference.
- 7. Click Save Security Settings.
- 8. Select **Confirm** and then click **Yes**. Security settings changed successfully.

Disable security key

About this task

(i) NOTE: Disabling Security Key is active if there is a security key present on the controller.

Steps

- 1. Enter the Dell PERC 11 Configuration Utility. See Navigate to Dell PERC 11 configuration utility.
- 2. Click Main Menu > Controller Management > Advanced Controller Management > Disable Security. You are prompted to confirm whether you want to continue.
- 3. Select the Confirm option.
- 4. Click Yes.
 - The security key is disabled successfully.

(i) NOTE: All virtual disks must be deleted or removed to disable security.



WARNING: Any un-configured secured disks in the system will be repurposed.

Create a secured virtual disk

About this task

To create a secured virtual disk, the controller must have a security key established first. See Create a security key. NOTE: Combining SAS and SATA hard drives within a virtual disk is not supported. Also, combining hard drives and olid-state drives (SSDs) within a virtual disk is not supported. Mixing of NVMe drives is not supported.

After the security key is established, perform the following steps:

Steps

- 1. Enter the Dell PERC 11 Configuration Utility. See Navigate to Dell PERC 11 configuration utility.
- Click Main Menu > Configuration Management > Create Virtual Disk. For more information, see Create virtual disks.
- 3. Select the Secure Virtual Disk option.
- 4. Click Create Virtual Disk. The secure virtual disk is created successfully.

Secure a non-RAID disk

In HII, secure a non-RAID disk by using the security key of the controller.

Steps

- 1. Enter the Dell PERC 11 Configuration Utility. See Navigate to Dell PERC 11 configuration utility.
- 2. Click Main Menu > Physical Disk Management.
- The list of Non-RAID disks is displayed.
- **3.** Select a non-RAID disk.
- 4. From the Operations drop-down menu, select Secure Non-RAID Disk.

Secure a pre-existing virtual disk

Steps

- 1. Enter the Dell PERC 11 Configuration Utility. See Navigate to Dell PERC 11 configuration utility.
- Click Main Menu > Virtual Disk Management. The list of virtual disks is displayed.
- 3. Select a virtual disk.
- 4. From the Operations drop-down menu, select Secure Virtual Disk.

(i) NOTE: The virtual disks can be secured only when the virtual disks are in Optimal state.

Import a secured non-RAID disk

If you are inserting a non-RAID disk into a system that has a controller key different from the security key on the drive, the security key from the system in which it was initially secured must be provided in HII.

Prerequisites

(i) NOTE: The controller must have an existing security key before importing a secured non-RAID disk.

84 Security key and RAID management



Steps

- 1. Enter the Dell PERC 11 Configuration Utility. See Navigate to Dell PERC 11 configuration utility.
- 2. Click Main Menu > Configuration Management > Manage Foreign Configurations.
- 3. Click Enter Passphrase for Locked Disks.
- A screen is displayed asking if you are sure you want to perform the operation.
- 4. Enter Passphrase if importing non-RAID disk with a different passphrase.
- 5. Select the **Confirm** option.
- 6. Click Yes.

(i) NOTE: If Auto-Configure for non-RAID Disks is enabled, the disk becomes a non-RAID disk. Else, it is unconfigured.

Import a secured virtual disk

Prerequisites

(i) NOTE: The controller must have an existing security key before importing secured foreign virtual disk.

Steps

- 1. Enter the Dell PERC 11 Configuration Utility. See Navigate to Dell PERC 11 configuration utility.
- 2. Click Main Menu > Configuration Management > Manage Foreign Configurations > Preview Foreign Configurations.
- 3. Click Import Foreign Configuration.
- A screen is displayed asking if you are sure you want to perform the operation.
- 4. Enter **Passphrase** if importing virtual disk with a different passphrase.
- 5. Select the Confirm option.
- 6. Click Yes.

The foreign configuration is imported successfully.

Dell Technologies OpenManage Secure Enterprise Key Manager

This feature allows the PERC to receive a security key from a remote server instead of saving the key on a local controller. This protects data on secured disks under the PERC if the disks or entire system is stolen. Refer to the www.dell.com/idracmanuals for more information on configuring OpenManage Secure Enterprise Key Manager, as well as Secure Sockets Layer (SSL) or Transport Layer Security (TLS) related configuration.

NOTE: Downgrade of PERC firmware to a firmware that does not support enterprise key management while enterprise key manager mode is enabled, is blocked.

- **NOTE:** When replacing a controller enabled with enterprise key management, lifecycle controller part replacement will re-configure the new controller to match the existing controller's configuration.
- () NOTE: If key exchange fails during boot, view and correct any connection issues with the key server identified in the iDRAC lifecycle log. Then the system can be cold booted.

Supported controllers for OpenManage Secure Enterprise Key Manager

Enterprise key manager mode is supported on the PERC H755 adapter, PERC H755 front SAS, and PERC H755N front NVMe, and allows the creation of secured virtual disks and non–RAID disks. For more information about supported platforms, see Support Site.



Enterprise key manager mode is not supported on the PERC H755 MX adapter, PERC H355 front SAS, PERC H355 adapter SAS, PERC H350 adapter SAS, and PERC H350 Mini Monolithic SAS.

Manage enterprise key manager mode

iDRAC manages Enterprise key manager features. For instructions on enabling enterprise key manager mode, vidi **dell.com/** idracmanuals.

- () NOTE: If preserved cache is present, the controller does not allow OpenManage Secure Enterprise Key Manager (SEKM) mode to be enabled.
- **NOTE:** When enterprise key manager mode is enabled, the controller waits up to two minutes for iDRAC to send keys, after which the PERC continues to boot.
- **NOTE:** Transitioning a controller from Local Key Management (LKM) mode to SEKM mode is supported on firmware starting with version 52.16.1-4074.
- **NOTE:** iDRAC performs rotation of keys. Any attempt to rekey the controller through a different management application is not supported.

Disable enterprise key manager mode

Enterprise key manager mode can be disabled from any supported Applications & User Interfaces supported by PERC 11. For more information, see the management application's user's guide or see Disable security key.

Manage virtual disks in enterprise key manager mode

Virtual disks are managed in the same way in enterprise key manager mode as in local key manager mode. SED capable virtual disks can be secured during or after creation. See Create a secured virtual disk.

Manage non-RAID disks in enterprise key manager mode

Non–RAID disks are managed in the same way in enterprise key manager mode as in local key manager mode. SED capable non–RAID disks can be secured after creation. See Create a secured virtual disk.

Transition of drives from local key management to enterprise key management (without supported firmware for PERC and iDRAC)

Local key management drives can be transitioned to an enterprise key management enabled system, but the controller cannot be transitioned from local key management mode to enterprise key manager mode or the reverse without first disabling security on the controller. Perform the following steps to transition from local key management drives to enterprise key management:

- 1. Save the current local key management security key.
- 2. Shut down both systems.
- 3. Remove the local key management drives and reinsert them to the enterprise key manager enabled system.
- 4. Power on the enterprise key manager system.
- 5. Go to HII foreign configuration.
- 6. Enter the local key management keys for those drives.
- 7. Import the configuration.
 - **NOTE:** Once local key management drives are migrated to enterprise key manager, they cannot be migrated back to local key management mode. The drives have to be cryptographically erased to disable security and then converted back to local key management disks. For more information about performing this action, contact Support Site.
- 86 Security key and RAID management



Migrate of drives from local key management to enterprise key management (with supported firmware for PERC and iDRAC)

PERC enables transition from Local Key Management (LKM) mode to Secure Enterprise Key Manager (SEKM) mode without disabling LKM security first. For instructions on transitioning from LKM mode to SEKM mode, see iDRAC Manuals.

(i) NOTE: This feature is supported on firmware starting with version 51.16.0-4076.

The transition from LKM to SEKM on the controller fails if the following are true at time of attempt:

- Snapdump is present on PERC.
- Preserved cache is present on PERC.
- RAID level migration is in progress on PERC,
- Online capacity expansion is in progress on PERC.
- Sanitize on a physical disk is in progress.
- LKM key that does not match with the current key of PERC.
- PERC firmware does not support transition.

Security key and RAID management 87



Troubleshooting issues in PERC11 cards

To get help for resolving issues in your PERC11 series cards, you can contact your Dell Technical Service representative. **Topics:**

- Single virtual disk performance or latency in hypervisor configurations
- Configured disks removed or not accessible error message
- Dirty cache data error message
- Discovery error message
- Drive Configuration Changes Error Message
- Windows operating system installation errors
- Firmware fault state error message
- Foreign configuration found error message
- Foreign configuration not found in HII
- Degraded state of virtual disks
- Memory errors
- Preserved Cache State
- Security key errors
- General issues
- Physical disk issues
- SMART errors
- Replace member errors
- Linux operating system errors
- Drive indicator codes
- HII error messages
- System reports more drive slots than what is available
- World Wide Number on drive sticker is not the same in applications
- Backplane firmware revision not changing in PERC interfaces after an update

Single virtual disk performance or latency in hypervisor configurations

Multi-initiator or hypervisor configurations running multiple I/O workloads to a single raid array may experience degraded performance or latency. This is caused by upper layers sending separate I/O workloads for each virtual machine to the storage subsystem which ends up being a random I/O workload to the under lying RAID array. For I/O workload configurations that require lower latency restrictions and higher I/O performance it may be beneficial to run fewer I/O workloads to individual RAID arrays or to use separate RAID arrays and physical disks for each I/O workload. Other considerations are making sure write-back, read ahead cache is enabled for rotational disks or using solid state drives (SSDs) to improve random I/O workload performance.

Performance degradation may also be observed when background operations such as initialization, consistency check, or reconstructions are running on the virtual disk. See your hypervisor storage best practices or performance best practices guides for additional configuration support.


Configured disks removed or not accessible error message

Error Message:	Some configured disks have been removed from your system or are no longer accessible. Check your cables and ensure all disks are present. Press any key or 'C' to continue.
Probable Cause:	The message indicates that some configured disks were removed. If the disks were not removed, they are no longer accessible. The cables from the PERC controller to the backplane might be improperly connected.
Corrective Action:	Check the cable connections and fix issues if any. Restart the system. If there are no cable problems, press any key or $<$ C $>$ to continue.

Dirty cache data error message

Error Message:	The following virtual disks are missing: (x). If you proceed (or load the configuration utility), these virtual disks will be removed from your configuration. If you wish to use them at a later time, they will have to be imported. If you believe these virtual disks should be present, please power off your system and check your cables to ensure all disks are present. Press any key to continue, or 'C' to load the configuration utility. The cache contains dirty data, but some virtual disks are missing
	or will go offline, so the cached data cannot be written to disk. If this is an unexpected error, then please power off your system and check your cables to ensure all disks are present. If you continue, the data in cache will be permanently discarded. Press 'X' to acknowledge and permanently destroy the cached data.
Probable Cause:	The controller preserves the dirty cache from a virtual disk if the disk becomes offline or is deleted because of missing physical disks. This message indicates that some configured disks were removed. If the disks were not removed, they are no longer accessible. The cables from the PERC controller to the backplane might be improperly connected.

CorrectiveCheck the cable connections and fix any problems. Restart the system. Use the HII configuration utility to
import the virtual disk or discard the preserved cache. For the steps to discard the preserved cache, see
Clear the cache memory.

Discovery error message

Error Message: A discovery error has occurred, please power cycle the system and all the enclosures attached to this system.
 Probable Cause: This message indicates that discovery did not complete within 120 seconds. The cables from the PERC controller to the backplane might be improperly connected.
 Corrective Action:

Drive Configuration Changes Error Message

Error Message: Entering the configuration utility in this state will result in drive configuration changes. Press 'Y' to continue loading the configuration utility or please power off your system and check your cables to ensure all disks are present and reboot.

Troubleshooting issues in PERC11 cards 89



Probable Cause:	The message is displayed after another HII warning indicating there are problems with previously configured disks and you have chosen to accept any changes and continue. The cables from the PERC controller to the backplane might be improperly connected.
Corrective	Check the cable connections and fix any problems before restarting the system. If there are no cable

Action:

Check the cable connections and fix any problems before restarting the system. If there are no cable problems, press any key or <Y> to continue.

Windows operating system installation errors

Ensure that you perform the following step before installing Windows on 4 KB sector drives:

1. Read and understand the updates to the version of Windows that you have installed. You can find this information in the Microsoft help. For more information, see Microsoft support policy for 4 K sector hard drives in Windows.

Firmware fault state error message

Error Message:Firmware is in Fault State.CorrectiveContact Global Technical Support.Action:

Foreign configuration found error message

Error Message:	Foreign configuration(s) found on adapter. Press any key to continue, or 'C' to load the configuration utility or 'F' to import foreign configuration(s) and continue.
Probable Cause:	When a controller firmware detects a physical disk with existing foreign metadata, it flags the physical disk as foreign and generates an alert indicating that a foreign disk was detected.
Corrective Action:	Press <f></f> at this prompt to import the configuration (if all member disks of the virtual disk are present) without loading the HII Configuration Utility . Or press <c></c> to enter the HII Configuration Utility and either import or clear the foreign configuration.

Foreign configuration not found in HII

Error Message:	The foreign configuration message is present during POST but no foreign configurations are present in the foreign view page in HII configuration utility. All virtual disks are in an optimal state.
Corrective	Ensure all your PDs are present and all VDs are in optimal state. Clear the foreign configuration using HII configuration utility or Dell OpenManage Server Administrator Storage Management.
Action:	CAUTION: The physical disk goes to Ready state when you clear the foreign configuration.

If you insert a physical disk that was previously a member of a virtual disk in the system, and that disk's previous location has been taken by a replacement disk through a rebuild, you must manually remove the foreign configuration flag of the newly inserted disk.

Degraded state of virtual disks

A redundant virtual disk is in a degraded state when one or more physical disks have failed or are inaccessible. For example, if a RAID 1 virtual disk consists of two physical disks and one of them fails or become inaccessible, the virtual disk become degraded.

To recover a virtual disk from a degraded state, you must replace the failed physical disk and rebuild it. Once the rebuilding process is complete, the virtual disk state changes from degraded to optimal.

90 Troubleshooting issues in PERC11 cards



Memory errors

Memory errors can corrupt cached data, so the controllers are designed to detect and attempt to recover from the memory errors. Single-bit memory errors can be handled by the controller and do not disrupt normal operation. A notification is sent if the number of single-bit errors exceeds a threshold value.

Multi-bit errors are more serious as they result in corrupted data and data loss. The following are the actions that occur in the case of multi-bit errors:

- If a multi-bit error occurs while accessing data in the cache when the controller is started with dirty cache, the controller discards the cache contents. The controller generates a warning message to the system console to indicate that the cache was discarded and generates an event.
- If a multi-bit error occurs at run-time either in code/data or in the cache, the controller stops.
- The controller logs an event to the controller's internal event log and a message during POST is displayed indicating a multi-bit error has occurred.

(i) NOTE: In case of a multi-bit error, contact Global Technical Support.

Preserved Cache State

The controller preserves the dirty cache from a virtual disk if the virtual disk goes offline or is deleted because of missing physical disks. This preserved dirty cache is called **pinned cache** and is preserved until you import the virtual disk or discard the cache.

- 1. Import the virtual disk—Power off the system, re-insert the virtual disk and restore the system power. Use the **HII Configuration Utility** to import the foreign configuration.
- 2. Discard the preserved cache—See Clear the cache memory.
- **NOTE:** It is recommended to clear the preserved cache before reboot using any of the virtual disks present on the controller.

Security key errors

Secured foreign import errors

A foreign configuration is a RAID configuration that already exists on a replacement physical disk that you install in a system. A secured foreign configuration is a RAID configuration that was created under a different security key.

There are two scenarios in which a secured foreign import fails:

- The passphrase authentication fails—A virtual disk secured with a security key different from the current controller security key cannot be imported without authentication of the original passphrase used to secure them. Supply the correct passphrase to import the secured foreign configuration. If you have lost or forgotten the passphrase, the secured foreign disks remain locked (inaccessible) until the appropriate passphrase is entered or if they are erased.
- The secured virtual disk is in an offline state after supplying the correct passphrase—You must check to determine why the virtual disk failed and correct the problem.

Failure to select or configure non Self-Encrypting Disks non-SED

A virtual disk can be either secured or unsecured depending on how it was configured when created. In order to create a secured virtual disk, the controller must have a security key present and must contain SEDs only. In order to select/configure non-SED, you must create an unsecured virtual disk. You can create an unsecured virtual disk even if there is a security key. Select the **Secure VD** option as **No** in the **Create New VD** menu. For steps on how to create an unsecured virtual disk, see Create virtual disks.



Failure to delete security key

A security key is used to lock or unlock access to a security-enabled component. This key is not utilized in the actual encryption of data. If a security key is present, both secured and unsecured virtual disks may exist.

To delete the security key, you must have a previously established security key present on the controller and there cannot be any configured secured disks. If there are configured secured virtual disks, remove or delete them.

Failure of Cryptographic Erase on encryption-capable physical disks

Cryptographic Erase is the process of securely erasing all data permanently on an encryption-capable physical disk and resetting the security attributes. It is used in scenarios such as deleting a foreign configuration in the event of a forgotten or lost passphrase or unlocking a disk that had been previously locked.

You can perform Cryptographic Erase only on encryption-capable disks that are not hot spares and not configured as non-RAID or virtual disks. Ensure that the conditions are met and see Cryptographic Erase.

General issues

PERC card has yellow bang in Windows operating system device manager

Issue:	The device is displayed in Device Manager but has a yellow bang (exclamation mark).
Corrective Action:	Reinstall the driver. For more information on reinstalling drivers, see Driver support for PERC 11.

PERC card not seen in operating systems

Issue:	The device does not appear in the Device Manager
Corrective	Power off the system and reseat the controller.

Issues in controller, battery, and disk when operating at low temperature

Issue:	If the controller is operating at temperatures less than zero degree Centigrade, then an increase in the number of issues related to controller, battery, or drive is observed.
Corrective Action:	Ensure that the controller ambient temperature is more than zero degree Centigrade.

Physical disk issues

Physical disk in failed state

Issue:

One of the physical disks in the disk array is in the failed state.

92 Troubleshooting issues in PERC11 cards



Corrective Update the PERC cards to the latest firmware available on the support site and replace the drive. **Action:**

Unable to rebuild a fault tolerant virtual disk

Issue:	Cannot rebuild a fault tolerant virtual disk. For more information, see the alert log for virtual disks.
Probable Cause:	The replacement disk is too small or not compatible with the virtual disk.
Corrective Action:	Replace the failed disk with a compatible good physical disk with equal or greater capacity.

Fatal error or data corruption reported

Issue:	$\label{eq:Fatalerror} \ensuremath{Fatal}\xspace \ensuremath{error}(s) \mbox{ or data corruption}(s) \mbox{ are reported when accessing virtual disks}.$
Corrective Action:	Contact Global Technical Support.

Multiple disks are inaccessible

Issue:	Multiple disks are simultaneously inaccessible.
Probable Cause:	Multiple physical disk errors in a single array typically indicate a failure in cabling or connection and could involve the loss of data.
Corrective Action:	You can recover the virtual disk after multiple physical disks become simultaneously inaccessible. Perform the following steps to recover the virtual disk:

CAUTION: Follow the safety precautions to prevent electrostatic discharge.

- 1. Turn off the system, check cable connections, and reseat physical disks.
- 2. Ensure that all the disks are present in the enclosure.
- 3. Turn on the system and enter the HII Configuration Utility.
- 4. Import the foreign configuration.
- 5. Press <F> at the prompt to import the configuration, or press <C> to enter the **HII Configuration Utility** and either import or clear the foreign configuration.

If the virtual disk is redundant and transitioned to **Degraded** state before going **Offline**, a rebuild operation starts automatically after the configuration is imported. If the virtual disk has gone directly to the **Offline** state due to a cable pull or power loss situation, the virtual disk is imported in its **Optimal** state without a rebuild occurring.

NOTE: You can use the **HII Configuration Utility** or Dell OpenManage storage management application to perform a manual rebuild of multiple physical disks.

Rebuilding data for a failed physical disk

Issue:	Rebuilding data for a physical disk that is in a failed state.
Probable Cause:	Physical disk is failed or removed.
Corrective Action:	If you have configured hot-spares, the PERC card automatically tries to use one of the hot-spares to rebuild a physical disk that is in a failed state. Manual rebuild is necessary if no hot-spares with enough capacity to rebuild the failed physical disks are available. You must insert a physical disk with enough storage in the subsystem before rebuilding the physical disk.

NOTE: You can use the **HII Configuration Utility** or Dell OpenManage storage management application to perform a manual rebuild of an individual physical disk.



Virtual disk fails during rebuild using a global hot spare

Issue:	A virtual disk fails during rebuild while using a global hot spare.
Probable Cause:	One or more disks in the virtual disks fails or is disconnected while the rebuild is in progress.
Corrective Action:	No action is required. The global hot spare reverts to Hot spare state and the virtual disk is in Failed state.

Dedicated hot spare disk fails during rebuild

Issue:	A hot spare disk fails during rebuild while using a dedicated hot spare.
Probable Cause:	The dedicated hot spare assigned to the virtual disk fails or is disconnected while the rebuild is in progress.
Corrective Action:	If there is a global hot spare available with enough capacity, rebuild will automatically start on the global hot spare. Where there is no hot spare present, you must insert a physical disk with enough capacity into the system before performing a rebuild.

Redundant virtual disk fails during reconstruction

Issue:	Multiple disks fails during a reconstruction process on a redundant virtual disk that has a hot spare.
Probable Cause:	Multiple physical disks in the virtual disk is failed or the cables are disconnected.
Corrective Action:	No action is required. The physical disk to which a reconstruction operation is targeted reverts to Ready state, and the virtual disk goes to Failed state. If there are any other virtual disks that can be supported by the capacity of the hot spare then the dedicated hot spare is converted to global hot spare, if not the hot spare will revert back to Ready state.

Virtual disk fails rebuild using a dedicated hot spare

Issue:	A virtual disk fails during rebuild while using a dedicated hot spare.
Probable Cause:	One or more disks in the virtual disks fails or is disconnected while the rebuild is in progress.
Corrective Action:	No action is required. The dedicated hot spare is in hot spare state and converted to global hot spare if there is any other virtual disk that is supported, otherwise the dedicated hot spare reverts to Ready state and the virtual drive is in Failed state.

Physical disk takes a long time to rebuild

Issue:	A physical disk is taking longer than expected to rebuild.
Description:	A physical disk takes longer to rebuild when under high I/O stress. There is only one rebuild I/O operation for every five host I/O operations.
Corrective Action:	If possible, reduce I/O stress on the physical disk or increase the value of rebuild rate controller parameter.

Drive removal and insertion in the same slot generates a foreign configuration event

Issue:

When a drive which is part of a virtual disk is removed and reinserted into the same slot the drive goes through a transient state of being foreign for a short period of time before rebuilding.

94 Troubleshooting issues in PERC11 cards



Description: This transient state could be reported as an event in management applications as A foreign configuration was detected on RAID Controller is SL x, where x is the slot of the RAID controller.

CorrectiveNo action is required on the foreign configuration state of the drive as it is transient and the controllerAction:handles the event automatically.

SMART errors

SMART monitors the internal performance of all motors, heads, and physical disk electronics and detects predictable physical disk failures.

NOTE: For information about SMART errors' reports that could indicate hardware failure, see the *Dell OpenManage Storage Management User's Guide* available at OpenManage Manuals.

Smart error detected on a non-RAID disk

Issue:

A SMART error is detected on a non-RAID disk.

Corrective Action:

- Perform the following steps: **1** Back up your data
 - 1. Back up your data.
 - 2. Replace the affected physical disk with a new physical disk of equal or higher capacity.
 - **3.** Restore from the backup.

Smart error detected on a physical disk in a non-redundant virtual disk

lssue:	A SMART error is detected on a physical disk in a non-redundant virtual disk.
Corrective	Perform the following steps:
Action:	1. Back up your data.
	2. Use Replace Member to replace the disk manually.

NOTE: For more information about the **Replace Member** feature, see Configure hot spare drives.

- 3. Replace the affected physical disk with a new physical disk of equal or higher capacity.
- **4.** Restore from the backup.

Smart error detected on a physical disk in a redundant virtual disk

lssue:	A SMART error is detected on a physical disk in a redundant virtual disk.
Corrective Action:	 Perform the following steps: 1. Back up your data. 2. Force the physical disk offline. () NOTE: If a hot spare is present, the rebuild starts with the hot spare after the disk is forced

offline.

- $\ensuremath{\textbf{3.}}$ Replace the disk with a new physical disk of equal or higher capacity.
- 4. Perform the **Replace Member** operation.

() NOTE: The **Replace Member** operation allows you to copy data from a source physical disk of a virtual disk to a target physical disk that is not a part of the virtual disk. For more information about the **Replace Member** feature, see the topic Configure hot spare drives.

Troubleshooting issues in PERC11 cards 95



Replace member errors

(i) NOTE: For more information about the **Replace Member** features, see Configure hot spare drives.

Source disk fails during replace member operation

lssue:	The source disk fails during the Replace Member operation and the Replace Member operation stops due to the source physical disk error.
Probable Cause:	Physical disk failure or physical disk is removed or disconnected.
Corrective Action:	No action required. If the virtual disk can tolerate disk failure, and the source data is available from other disks in the virtual disk, the rebuild begins automatically on the target disk, using the data from the other disks, if the virtual disk cannot tolerate the failure, the virtual disk goes to offline state and the replace member operation is stopped.

Target disk fails during replace member operation

Issue:	The target disk failure reported during the Replace Member operation, and the Replace Member operation stops.
Probable Cause:	Physical disk failure or physical disk is removed or disconnected.
Corrective Action:	It is recommended that you replace or check the target drive, and restart the Replace Member operation or perform the operation on a different target drive.

A member disk failure is reported in the virtual disk which undergoes replace member operation

Issue:The source and the target drive which is part of Replace Member operation are online, while a different
drive which is a member of the virtual drive reports a failure.Probable Cause:Physical disk failure or physical disk is removed or disconnected.Corrective
Action:A rebuild starts if there any hot-spares configured or you may replace the failed drive. The Replace
Member operation continues as far as the source virtual disk can tolerate the drive failure. If the source
virtual disk fails, the Replace Member is stopped, otherwise the virtual disk continues to be in degraded
state

Linux operating system errors

Virtual disk policy is assumed as write-through

Error:<Date:Time> <HostName> kernel: sdb: asking for cache data failed<Date:Time> <HostName> kernel: sdb: assuming drive cache: write throughCorrective
Action:The error message is displayed when the Linux Small Computer System Interface (SCSI) mid-layer asks for physical disk cache settings. The controller firmware manages the virtual disk cache settings on a per controller and a per virtual disk basis, so the firmware does not respond to this command. The Linux SCSI mid-layer assumes that the virtual disk's cache policy is Write-Through. SDB is the device node for a virtual disk. This value changes for each virtual disk. Except for this message, there is no effect of this behavior on normal operation. The cache policy of the virtual disk and the I/O throughput are not affected by this message. The cache policy settings for the PERC SAS RAID system remain unchanged.



Unable to register SCSI device error message

Error:

smartd[smartd[2338] Device: /dev/sda, Bad IEC (SMART) mode page, err=-5, skip device smartd[2338] Unable to register SCSI device /dev/sda at line 1 of file /etc/smartd.conf.

CorrectiveThis is a known issue. An unsupported command is entered through the user application. User applications
attempt to direct Command Descriptor Blocks to RAID volumes. The error message does not affect the
feature functionality. The Mode Sense/Select command is supported by firmware on the controller.
However, the Linux kernel daemon issues the command to the virtual disk instead of to the driver IOCTL
node. This action is not supported.

Drive indicator codes

The LEDs on the drive carrier indicates the state of each drive. Each drive carrier has two LEDs: an activity LED (green) and a status LED (bicolor, green/amber). The activity LED blinks whenever the drive is accessed.



Figure 30. Drive indicators

- 1. Drive activity LED indicator
- 2. Drive status LED indicator
- 3. Drive capacity label

If the drive is in the Advanced Host Controller Interface (AHCI) mode, the status LED indicator does not power on. Drive status indicator behavior is managed by Storage Spaces Direct. Not all drive status indicators may be used.

Table 21. Drive indicator codes

Drive status indicator code	Condition
Blinks green twice per second	The drive is being identified or preparing for removal
Off	The drive is ready for removal () NOTE: The drive status indicator remains off until all drives are initialized after the system is powered on. Drives are not ready for removal during this time.
Blinks green, amber, and then powers off	There is an expected drive failure
Blinks amber four times per second	The drive has failed
Blinks green slowly	The drive is rebuilding
Solid green	The drive is online
Blinks green for three seconds, amber for three seconds, and then powers off after six seconds	The rebuild has stopped



HII error messages

Unhealthy Status of the drivers

Error:	One or more boot driver(s) have reported issues. Check the Driver Health Menu in Boot Manager for details.
Probable Cause:	This message might indicate that the cables are not connected, the disks might be missing, or the UEFI driver might require configuration changes.
Corrective Action:	 Check if the cables are connected properly, or replace missing hard drives, if any and then restart the system. Press any key to load the driver health manager to display the configurations. The Driver Health Manager displays the driver(s), which requires configuration.
	3. Alternately, if the UEFI driver requires configuration, press any key to load the Configuration Utility.

Rebuilding a drive during full initialization

Issue:	Automatic rebuild of drives is disabled for virtual disk during full initialization.
Corrective Action:	After full initialization the drive will automatically start its rebuild on its corresponding virtual disk.

System reports more drive slots than what is available

The system reports more slots than what is available in the following two scenarios:

System drives are hot swappable with backplane.	When the system drives are hot swappable, the PERC controller is not able to communicate correctly with the backplane or enclosure. Hence, the PERC controller reports a generic enclosure with drive 16 slots. In iDRAC, under Overview > Enclosures , the Enclosure ID is displayed as BP_PSV and Firmware version is displayed as 03 .
Corrective action	Turn off the system, reseat the controller and all the cables on the controller and backplane. If the issue is not resolved, contact your Dell Technical Service representative.
System drives are not hot swappable with cable direct attached.	When the system drives are not hot swappable, a default enclosure with 16 drive slots is expected to be reported (even though the system does not support that many drives).

World Wide Number on drive sticker is not the same in applications

World Wide Number (WWN) on the drive sticker and applications are not matching. NVMe drives do not have a WWN. So, the applications create a WWN from the available drive information. This WWN may not match with the WWN that is on the drive sticker, if present.

98 Troubleshooting issues in PERC11 cards



Backplane firmware revision not changing in PERC interfaces after an update

After updating the backplane firmware on 15G and later PowerEdge servers, the backplane version will not show as updated on some interfaces until the system is reset.

Troubleshooting issues in PERC11 cards 99



Appendix RAID description

RAID is a group of independent physical disks that provides high performance by increasing the number of disks used for saving and accessing data.

CAUTION: In the event of a physical disk failure, a RAID 0 virtual disk fails, resulting in data loss.

A RAID disk subsystem offers the following benefits:

- Improved I/O performance and data availability.
- Improved data throughput because several disks are accessed simultaneously. The physical disk group appears either as a single storage unit or multiple logical units to the host system.
- Improved data storage availability and fault tolerance. Data loss caused by a physical disk failure can be recovered by rebuilding missing data from the remaining physical disks containing data or parity.

Topics:

- Summary of RAID levels
- RAID 10 configuration
- RAID terminology

Summary of RAID levels

Following is a list of the RAID levels supported by the PERC 11 series of cards:

- RAID 0 uses disk striping to provide high data throughput, especially for large files in an environment that requires no data redundancy.
- RAID 1 uses disk mirroring so that data written to one physical disk is simultaneously written to another physical disk. RAID 1 is good for small databases or other applications that require small capacity and complete data redundancy.
- RAID 5 uses disk striping and parity data across all physical disks (distributed parity) to provide high data throughput and data redundancy, especially for small random access.
- RAID 6 is an extension of RAID 5 and uses an additional parity block. RAID 6 uses block-level striping with two parity blocks distributed across all member disks. RAID 6 provides protection against double disk failures, and failures while a single disk is rebuilding. If you are using only one array, deploying RAID 6 is more effective than deploying a hot spare disk.
- RAID 10 is a combination of RAID 0 and RAID 1, uses disk striping across mirrored disks. It provides high data throughput and complete data redundancy.
- RAID 50 is a combination of RAID 0 and RAID 5 where a RAID 0 array is striped across RAID 5 elements. RAID 50 requires at least six disks.
- RAID 60 is a combination of RAID 0 and RAID 6 where a RAID 0 array is striped across RAID 6 elements. RAID 60 requires at least eight disks.

The following table lists the minimum and maximum disks supported on each RAID levels.

Table 22. Minimum and maximum disks supported on each RAID levels

RAID Level	Minimum disk	Maximum disk
0	1	32
1	2	2
5	3	32
6	4	32
10	4	240
50	6	240
60	8	240



(i) NOTE: The maximum number of virtual disks is currently limited to 192, because of the supported enclosure configuration.

RAID 10 configuration

In PERC 10 and PERC 11 controllers, RAID 10 can be configured without spanning up to 32 drives. Any RAID 10 volume that has more than 32 drives require spanning. Each span can contain up to 32 drives. Drives must be distributed evenly across all the spans with each span containing an even number of drives.

NOTE: Spans in a RAID 10 volume are only supported if spans are even. Uneven spanned RAID 10 cannot be imported from previous controller generations.

The following table shows the RAID 10 configurations.

Table 23. RAID 10 configurations

Disk or span count	RAID 10 capable						
4 (1)	Yes	64 (2)	Yes	124	No	184	No
6 (1)	Yes	66 (3)	Yes	126 (7)	Yes	186	No
8 (1)	Yes	68	No	128 (4)	Yes	188	No
10 (1)	Yes	70 (5)	Yes	130 (5)	Yes	190	No
12 (1)	Yes	72 (3)	Yes	132 (6)	Yes	192 (6)	Yes
14 (1)	Yes	74	No	134	No	194	No
16 (1)	Yes	76	No	136	No	196 (7)	Yes
18 (1)	Yes	78 (3)	Yes	138	No	198	No
20 (1)	Yes	80 (4)	Yes	140 (5)	Yes	200	No
22 (1)	Yes	82	No	142	No	202	No
24 (1)	Yes	84 (6)	Yes	144	Yes	204	No
26 (1)	Yes	86	No	146	No	206	No
28 (1)	Yes	88 (4)	Yes	148	No	208 (8)	Yes
30 (1)	Yes	90 (3)	Yes	150 (5)	Yes	210 (7)	Yes
32 (1)	Yes	92	No	152	No	212	No
34	No	94	No	154 (7)	Yes	214	No
36 (2)	Yes	96 (3)	Yes	156 (6)	Yes	216	No
38	No	98 (7)	Yes	158	No	218	No
40 (2)	Yes	100 (5)	Yes	160 (5)	Yes	220	No
42 (2)	Yes	102	No	162	No	222	No
44 (2)	Yes	104 (4)	Yes	164	No	224 (8)	Yes
46	No	106	No	166	No	226	No
48 (2)	Yes	108 (6)	Yes	168 (6)	Yes	228	No
50 (2)	Yes	110 (5)	Yes	170	No	230	No
52 (2)	Yes	112 (4)	Yes	172	No	232	No
54 (2)	Yes	114	No	174	No	234	No
56 (2)	Yes	116	No	176 (8)	Yes	236	No
58	No	118	No	178	No	238	No

Appendix RAID description 101



Table 23. RAID 10 configurations (continued)

Disk or span count	RAID 10 capable						
60 (2)	Yes	120 (4)	Yes	180 (6)	Yes	240 (8)	Yes
62	No	122	No	182 (7)	Yes	-	-

RAID terminology

Disk striping

Disk striping allows you to write data across multiple physical disks instead of just one physical disk. Disk striping involves partitioning each physical disk storage space in stripes of the following sizes: 64 KB, 128 KB, 256 KB, 512 KB, and 1 MB. The stripes are interleaved in a repeated sequential manner. The part of the stripe on a single physical disk is called a stripe element.

For example, in a four-disk system using only disk striping (used in RAID 0), segment 1 is written to disk 1, segment 2 is written to disk 2, and so on. Disk striping enhances performance because multiple physical disks are accessed simultaneously, but disk striping does not provide data redundancy.



Figure 31. Example of disk striping (RAID 0)

Disk mirroring

With mirroring (used in RAID 1), data written to one disk is simultaneously written to another disk. If one disk fails, the contents of the other disk can be used to run the system and rebuild the failed physical disk. The primary advantage of disk mirroring is that it provides complete data redundancy. Both disks contain the same data at all times. Either of the physical disks can act as the operational physical disk.

Disk mirroring provides complete redundancy, but is an expensive option because each physical disk in the system must be duplicated.





Stripe element 2 Stripe element 2 Duplicated Stripe element 3 Stripe element 3 Duplicated Stripe element 3 Stripe element 4 Duplicated

Figure 32. Example of Disk Mirroring (RAID 1)



Spanned RAID levels

Spanning is a term used to describe the way in which RAID levels 10, 50, and 60 are constructed from multiple sets of basic, or simple RAID levels. For example, a RAID 10 has multiple sets of RAID 1 arrays where each RAID 1 set is considered a span. Data is then striped (RAID 0) across the RAID 1 spans to create a RAID 10 virtual disk. Similarly, RAID 50 and RAID 60 combine multiple sets of RAID 5 or RAID 6 respectively with striping.

Parity data

Parity data is redundant data that is generated to provide fault tolerance within certain RAID levels. In the event of a disk failure, the parity data can be used by the controller to regenerate user data. Parity data is present for RAID 5, 6, 50, and 60.

The parity data is distributed across all the physical disks in the system. If a single physical disk fails, it can be rebuilt from the parity and the data on the remaining physical disks. RAID level 5 combines distributed parity with disk striping. Parity provides redundancy for one physical disk failure without duplicating the contents of the entire physical disks.

RAID 6 combines dual distributed parity with disk striping. This level of parity allows for two disk failures without duplicating the contents of entire physical disks.



Shipe element 1	Stripe element 2	Since element 3.	Stipe element 4	Shipe element 5	Parity (1-6)
Shipe element 7	Stripe element 8	Stripe element 9	Stripe element 10	Parity (6-10)	Strips element 6
Stripe element 13	Stripe element 14	Strips element 15	Parity (11-15)	Stripe element 11	Shipe element 12
Stripe element 19	Stripe element 20	Parity (16-20)	Stripe element 16	Shipe element 17	Stripe element 18
Stripe element 25	Parity (21-25)	Stripe element 21	Stripe element 22	Siripe element 23	Stripe element 24
Parity (26-30)	Stripe element 28	Stripe element 27	Stripe element 28	Stripe-element 29	Stripe element 30

Figure 33. Example of Distributed Parity (RAID 5)

(i) NOTE: Parity is distributed across multiple physical disks in the disk group.



Figure 34. Example of Dual Distributed Parity (RAID 6)

(i) NOTE: Parity is distributed across all disks in the array.



Getting help

Topics:

- Recycling or End-of-Life service information
- Contacting Dell
- Locating the Express Service Code and Service Tag
- Receiving automated support with SupportAssist

Recycling or End-of-Life service information

Take back and recycling services are offered for this product in certain countries. If you want to dispose of system components, visit the How to Recycle page and select the relevant country.

Contacting Dell

Dell provides online and telephone based support and service options. If you do not have an active internet connection, you can find Dell contact information on your purchase invoice, packing slip, bill or Dell product catalog. The availability of services varies depending on the country and product, and some services may not be available in your area. To contact Dell for sales, technical assistance, or customer service issues:

Steps

- 1. Go to the Support site.
- 2. Select your country from the drop-down menu on the lower right corner of the page.
- **3.** For customized support:
 - a. Enter the system Service Tag in the Enter a Service Tag, Serial Number, Service Request, Model, or Keyword field.
 - b. Click Submit.

The support page that lists the various support categories is displayed.

- 4. For general support:
 - **a.** Select your product category.
 - b. Select your product segment.
 - c. Select your product.
 - The support page that lists the various support categories is displayed.
- **5.** For contact details of Dell Global Technical Support:
 - a. Click Global Technical Support.
 - b. The **Contact Technical Support** page is displayed with details to call, chat, or e-mail the Dell Global Technical Support team.

Locating the Express Service Code and Service Tag

The unique Express Service Code and Service Tag is used to identify the system.

The information tag is located on the front of the system rear of the system that includes system information such as Service Tag, Express Service Code, Manufacture date, NIC, MAC address, QRL label, and so on. If you have opted for the secure default access to iDRAC, the Information tag also contains the iDRAC secure default password. If you have opted for iDRAC Quick Sync 2, the Information tag also contains the OpenManage Mobile (OMM) label, where administrators can configure, monitor, and troubleshoot the PowerEdge servers.





Figure 35. Locating the Express Service Code and Service tag

- 1. Information tag (front view)
- 3. OpenManage Mobile (OMM) label
- 5. Service Tag, Express Service Code, QRL label
- 2. Information tag (back view)
- 4. iDRAC MAC address and iDRAC secure password label

The Mini Enterprise Service Tag (MEST) label is located on the rear of the system that includes Service Tag (ST), Express Service Code (Exp Svc Code), and Manufacture Date (Mfg. Date). The Exp Svc Code is used by Dell to route support calls to the appropriate personnel.

Alternatively, the Service Tag information is located on a label on left wall of the chassis.

Receiving automated support with SupportAssist

Dell SupportAssist is an optional Dell Services offering that automates technical support for your Dell server, storage, and networking devices. By installing and setting up a SupportAssist application in your IT environment, you can receive the following benefits:

- Automated issue detection SupportAssist monitors your Dell devices and automatically detects hardware issues, both proactively and predictively.
- Automated case creation When an issue is detected, SupportAssist automatically opens a support case with Dell Technical Support.
- Automated diagnostic collection SupportAssist automatically collects system state information from your devices and uploads it securely to Dell. This information is used by Dell Technical Support to troubleshoot the issue.
- Proactive contact A Dell Technical Support agent contacts you about the support case and helps you resolve the issue.

The available benefits vary depending on the Dell Service entitlement purchased for your device. For more information about SupportAssist, go to the SupportAssist page.



Documentation resources

This section provides information about the documentation resources for your system.

To view the document that is listed in the documentation resources table:

- From the Dell support site:
 - 1. Click the documentation link that is provided in the Location column in the table.
 - 2. Click the required product or product version.

(i) NOTE: To locate the product name and model, see the front of your system.

3. On the Product Support page, click Manuals & documents.

- Using search engines:
 - Type the name and version of the document in the search box.

Table 24. Additional documentation resources for your system

Task	Document	Location
Setting up your system	For more information about installing and securing the system into a rack, see the Rail Installation Guide included with your rail solution. For information about setting up your system, see the <i>Getting Started Guide</i> document that is shipped with your system.	PowerEdge Server Manuals
Configuring your system	For information about the iDRAC features, configuring and logging in to iDRAC, and managing your system remotely, see the Integrated Dell Remote Access Controller User's Guide.	PowerEdge Server Manuals
	For information about understanding Remote Access Controller Admin (RACADM) subcommands and supported RACADM interfaces, see the RACADM CLI Guide for iDRAC.	
	For information about Redfish and its protocol, supported schema, and Redfish Eventing implemented in iDRAC, see the Redfish API Guide.	
	For information about iDRAC property database group and object descriptions, see the Attribute Registry Guide.	
	For information about Intel QuickAssist Technology, see the Integrated Dell Remote Access Controller User's Guide.	
	For information about earlier versions of the iDRAC documents.	iDRAC Manuals
	To identify the version of iDRAC available on your system, on the iDRAC web interface, click	



Table 24. Additional documentation resources for your system (continued)

Task	Document	Location
	? > About.	
	For information about updating drivers and firmware, see the Methods to download firmware and drivers section in this document.	Drivers
Understanding event and error messages	For information about the event and error messages generated by the system firmware and agents that monitor system components, go to qrl.dell.com > Look Up > Error Code, type the error code, and then click Look it up.	PowerEdge Server Event and Error Messages

Documentation resources 107



Home > Hardware > PowerEdge R760XS

By Dell Inc

PowerEdge R760XS

No description provided.

Product type	Hardware
Processor architecture	X64
Vendor	Dell Inc

is a certification that ensures hardware, like servers and storage devices, meets specific reliability and quality standards. Read more about <u>Additional Qualification</u> <u>Certification</u>.

Vindows	Product type	Hardware
erver 025	Manufacturer	Dell Inc.
a stiff a st	Processor name	INTEL(R) XEON(R) GOLD 6534
Lertified	Tested memory	1536GB 5600MHz

 \mathcal{O}

- > View Additional Qualifications and features
- > View and download submission(s)

Windows	Product type	Hardware
Server 2022	Manufacturer	Dell Inc.
	Processor name	INTEL(R) XEON(R) GOLD 6534
Certified	Tested memory	1536GB 5600MHz



D&LLTechnologies

Specification Sheet





PowerEdge R760xs

Best choice in balanced compute and flexible storage for the most popular IT applications

Buy the performance and flexibility you need

The new Dell PowerEdge R760xs is a 2U, two-socket rack server. Buy the best fit in scalable performance and large storage capability with this purpose-built 2U system. Focused on delivering the latest technology to power the most popular applications and workloads used by businesses today, including virtual desktop infrastructure (VDI), virtual machines (VMs), and software-defined storage (SDS). All delivered in a thoughtfully crafted platform that will provide balanced compute that fits in your current infrastructure.

Easily configurable

- Add up to two 4th generation Intel® Xeon® Scalable processors with up to 32 cores per socket for faster performance
- · Accelerate in-memory workloads with up to 16 DDR5 RDIMMS up to 4800 MT/sec
- Improve data throughput and reduce latency with support up to 8 I/O device (up to available 6 PCIe slots, 1 OCP 3.0 networking slot, and 1 dedicated PERC slot)
- Storage options include up to 12x 3.5" HDDs/SSDs, or up to 16x 2.5" HDD/SSDs, plus up to 8x NVMe drives

A breeze to cool

- Thoughtfully designed to fit in your current air-cooled infrastructure
- · Alleviate the worry about expensive liquid cooling retrofitting to your data center
- Synchronize your workload needs with a tailored performance configuration that is air cooled
- Minimize the carbon footprint of your data center by better matching the system power consumption with anticipated workload requirements

Cyber Resilient Architecture for Zero Trust IT environment & operations

Security is integrated into every phase of the PowerEdge lifecycle, including protected supply chain and factory-to-site integrity assurance. Silicon-based root of trust anchors end-to-end boot resilience while Multi-Factor Authentication (MFA) and role-based access controls ensure trusted operations.

Increase efficiency and accelerate operations with autonomous collaboration

The Dell OpenManage[™] systems management portfolio delivers a secure, efficient, and comprehensive solution for PowerEdge servers. Simplify, automate and centralize one-to-many management with the OpenManage Enterprise console and iDRAC.

Sustainability

From recycled materials in our products and packaging, to thoughtful, innovative options for energy efficiency, the PowerEdge portfolio is designed to make, deliver, and recycle products to help reduce the carbon footprint and lower your operation costs. We even make it easy to retire legacy systems responsibly with Dell Technologies Services.

Rest easier with Dell Technologies Services

Maximize your PowerEdge Servers with comprehensive services ranging from Consulting, to ProDeploy and ProSupport suites, Data Migration and more – available across 170 locations and backed by our 60K+employees and partners.

PowerEdge R760xs

The Dell PowerEdge R760xs offers compelling performance in a right-sized system with the latest PCIe Gen 5 bandwidth and large storage capability to support:

- Virtual Desktop Infrastructure (VDI)
- Virtual Machines (VMs)
- Software-Defined Storage Node

Feature	Technical Specifications	
Processor	Up to two 4th Generation Intel Xeon Scalable processor with up to 32 of	cores per processor
Memory	• 16 DDR5 DIMM slots, supports RDIMM 1 TB max, speeds up to 4	800 MT/s, supports registered ECC DDR5 DIMMs only
Storage controllers	 Internal Controllers: PERC H965i, PERC H755, PERC H755N, PE Internal Boot: Boot Optimized Storage Subsystem (BOSS-N1): HV External HBA (non-RAID): HBA355e; Software RAID: S160 	:RC H355, HBA355i NRAID 1, 2 x M.2 NVMe SSDs or USB
GPU Options	2 x 75 W SW, LP	
Drive Bays	 Front bays: 0 drive bay Up to 8 x 3.5-inch SAS/SATA (HDD/SSD) max 160 TB Up to 12 x 3.5-inch SAS/SATA (HDD/SSD) max 240 TB Up to 8 x 2.5-inch SAS/SATA/NVMe (HDD/SSD) max 122.88 TB Up to 16 x 2.5-inch SAS/SATA (HDD/SSD) max 121.6 TB Up to 16 x 2.5-inch (SAS/SATA) + 8 x 2.5-inch (NVMe) (HDD/SSD) max 244.48 TB 	 Rear bays: Up to 2 x 2.5-inch SAS/SATA/NVMe (HDD/SSD) max 30.72 TB (supported only with 12 x 3.5-inch SAS/SATA HDD/SSD configuration)
Hot swap Redundant Power Supplies	 1800 W Titanium 200—240 VAC or 240 HVDC 1400 W Platinum 100—240 VAC or 240 HVDC 1100 W Titanium 100—240 VAC or 240 HVDC 1100 W LVDC -48 — (-60) VDC 	 800 W Platinum 100—240 VAC or 240 HVDC, 700 W Titanium 200—240 VAC or 240 HVDC 600 W Platinum 100—240 VAC or 240 HVDC
Cooling Options	Air cooling	
Fans Dimensions and Weight	 Standard (STD) fans/High performance Silver (HPR) fans/ High performance Silver (HPR) fans/ H	 erformance Gold (VHP) fans, Up to 6 hot swappable fans Depth – 707.78 mm (27.85 inches) – without bezel 721.62 mm (28.4 inches) – with bezel Weight – Max 28.6 kg (63.0 lbs.)
Form Factor	2U rack server	J J J J J J J J J J J J J J J J J J J
Embedded Management	 iDRAC9 iDRAC Direct iDRAC RESTful API with Redfish 	iDRAC Service ModuleQuick Sync 2 wireless module
Bezel	Optional LCD bezel or security bezel	
OpenManage Software	 CloudIQ for PowerEdge plug in OpenManage Enterprise OpenManage Enterprise Integration for VMware vCenter OpenManage Integration for Microsoft System Center 	 OpenManage Integration with Windows Admin Center OpenManage Power Manager plugin OpenManage Service plugin OpenManage Update Manager plugin
Mobility	OpenManage Mobile	
OpenManage Integrations	 BMC Truesight Microsoft System Center OpenManage Integration with ServiceNow 	 Red Hat Ansible Modules Terraform Providers VMware vCenter and vRealize Operations Manager
Security	 Cryptographically signed firmware Data at Rest Encryption (SEDs with local or external key mgmt) Secure Boot Secure Erase 	 Secured Component Verification (Hardware integrity check) Silicon Root of Trust System Lockdown (requires iDRAC9 Enterprise or Datacenter) TPM 2.0 FIPS, CC-TCG certified, TPM 2.0 China NationZ
Embedded NIC	2 x 1 GbE LOM	
Network options	1 x OCP card 3.0 (optional)	
Ports	Front Ports: • 1 x iDRAC Direct (Micro-AB USB) port, 1 x USB 2.0, 1 x VGA Internal Ports: 1 x USB 3.0 (optional)	 Rear Ports 1 x Dedicated iDRAC Ethernet port, 1 x USB 2.0, 1 x USB 3.0, 1 x VGA, 1 x Serial (optional)
PCIe	 1 CPU Configuration: Up to 4 PCIe slots (2 x8 Gen5, 1 x16 Gen4, 2 CPU configuration: Up to 6 PCIe slots (2 x16 Gen5, 3 x16 Gen4 	1 x8 Gen4) , 1 x8 Gen4)
Operating System and Hypervisors	Microsoft Windows Server with Hyper-V Red Hat Enterprise Linux SUSE Linux Enterprise Server	VMware ESXi Canonical Ubuntu Server LTS For specifications and interoperability details, see Dell com/OSsupport
OEM-ready version available	From bezel to BIOS to packaging, your servers can look and feel as if solutions -> OEM Solutions.	they were designed and built by you. For more information, visit Dell.com ->

APEX Flex on Demand

Acquire the technology you need to support your changing business with payments that scale to match actual usage. For more information, visit https://www.delltechnologies.com/en-us/payment-solutions/flexible-consumption/flex-on-demand.htm.



Copyright © 2023 Dell Inc. or its subsidiaries. All Rights Reserved. Dell Technologies, Dell, and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be trademarks of their respective owners.

D&LTechnologies

April 2023



Dell PowerEdge R760xs

Technical Guide

Regulatory Model: E88S Regulatory Type: E88S001 November 2024 Rev. A03





Notes, cautions, and warnings

(i) NOTE: A NOTE indicates important information that helps you make better use of your product.

CAUTION: A CAUTION indicates either potential damage to hardware or loss of data and tells you how to avoid the problem.

WARNING: A WARNING indicates a potential for property damage, personal injury, or death.

© 2023 Dell Inc. or its subsidiaries. All rights reserved. Dell Technologies, Dell, and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be trademarks of their respective owners.



Chapter 1: System overview	5
Key workloads	5
New technologies	5
Chapter 2: System features and generational comparison	7
Chapter 3: Chassis views and features	9
Chassis views	
Front view of the system	9
Rear view of the system	
Inside the system	
System diagnostics and indicator codes	
QR code for PowerEdge R760xs system resources	18
Chapter 4: Processor	19
Processor features	
Supported processors	19
Chipset	
, Chipset features	20
Chapter 5: Memory subsystem	21
Supported memory	
General memory module installation guidelines	21
General memory module installation guidelines	21 23
General memory module installation guidelines Chapter 6: Storage Storage controllers	
General memory module installation guidelines Chapter 6: Storage Storage controllers	
General memory module installation guidelines Chapter 6: Storage Storage controllers Storage controller feature matrix Server storage controllers User Guide	
General memory module installation guidelines Chapter 6: Storage Storage controllers Storage controller feature matrix Server storage controllers User Guide RAID - Redundant Array of Independent Disks.	
General memory module installation guidelines Chapter 6: Storage Storage controllers Storage controller feature matrix Server storage controllers User Guide RAID - Redundant Array of Independent Disks Datasheets and PERC performance scaling decks.	
General memory module installation guidelines Chapter 6: Storage Storage controllers Storage controller feature matrix Server storage controllers User Guide RAID - Redundant Array of Independent Disks Datasheets and PERC performance scaling decks Boot Optimized Storage Solution (BOSS).	
General memory module installation guidelines Chapter 6: Storage Storage controllers Storage controller feature matrix Server storage controllers User Guide RAID - Redundant Array of Independent Disks Datasheets and PERC performance scaling decks Boot Optimized Storage Solution (BOSS) Supported Drives.	
General memory module installation guidelines Chapter 6: Storage Storage controllers Storage controller feature matrix Server storage controllers User Guide RAID - Redundant Array of Independent Disks Datasheets and PERC performance scaling decks Boot Optimized Storage Solution (BOSS) Supported Drives External Storage	21 23 23 23 25 25 25 25 25 25 27 27
General memory module installation guidelines Chapter 6: Storage Storage controllers Storage controller feature matrix Server storage controllers User Guide RAID - Redundant Array of Independent Disks Datasheets and PERC performance scaling decks Boot Optimized Storage Solution (BOSS) Supported Drives External Storage	21 23 23 23 25 25 25 25 25 25 27 27
General memory module installation guidelines Chapter 6: Storage Storage controllers Storage controller feature matrix Server storage controllers User Guide RAID - Redundant Array of Independent Disks Datasheets and PERC performance scaling decks Boot Optimized Storage Solution (BOSS) Supported Drives External Storage	
General memory module installation guidelines Chapter 6: Storage Storage controllers	21 23 23 25 25 25 25 25 25 27 27 27 27 27 27 27
General memory module installation guidelines Chapter 6: Storage Storage controllers	
General memory module installation guidelines	
General memory module installation guidelines Chapter 6: Storage Storage controllers Storage controller feature matrix Server storage controllers User Guide RAID - Redundant Array of Independent Disks Datasheets and PERC performance scaling decks Boot Optimized Storage Solution (BOSS) Supported Drives External Storage Chapter 7: Networking Overview. OCP 3.0 support. Supported OCP cards OCP NIC 3.0 vs. rack Network Daughter Card comparisons	
General memory module installation guidelines Chapter 6: Storage Storage controllers	
General memory module installation guidelines Chapter 6: Storage Storage controllers Storage controller feature matrix Server storage controllers User Guide RAID - Redundant Array of Independent Disks Datasheets and PERC performance scaling decks Boot Optimized Storage Solution (BOSS) Supported Drives External Storage Chapter 7: Networking OVERVIEW	
General memory module installation guidelines	



Chapter 9: Power, thermal, and acoustics	
Power	
PSU specifications	
Thermal	41
Acoustics	41
Acoustical configurations of R760xs	41
Chapter 10: Rack, rails, and cable management	
Rails information	43
A11 Sliding Rails features summary	43
A8 Static Rails features summary	
Cable Management Arm	
Strain Relief Bar	
Rack Installation	
Chapter 11: Supported operating systems	51
Chapter 12: Dell OpenManage Systems Management	52
Integrated Dell Remote Access Controller (iDRAC)	
Systems Management software support matrix	
Chapter 13: Appendix A. Standards compliance	55
Chapter 14: Appendix B: Additional specifications	
Chassis dimensions	
System weight	
Video specifications.	
USB ports specifications.	
Environmental specifications	86
I hermal restriction matrix	
Chapter 15: Appendix C Additional resources	64
	6b
Customer kits	
Dell Upgrades	
Upgrades portfolio	
Upgrades reference links	
Chapter 16: Appendix D: Service and support	
Why attach service contracts	
ProSupport Intrastructure Suite	
Speciality Support Services	/U
Dell deployment services	
Supplemental Deployment Services	/b
Unique Deployment Scenarios.	
DAY Z - AUTOMATION Services WITH ANSIDE	
Deil Technologies Consulting Services	



System overview

The Dell PowerEdge R760xs is Dell's latest two-socket, rack server that is designed to run complex workloads using highly scalable memory, I/O, and network options. The systems feature the 4thand 5th Gen Intel® Xeon® Scalable Processor (Socket E1/LGA4677-1), up to 16 DIMMs, PCI Express® (PCIe) 5.0 enabled expansion slots , and a broad selection of network interface technologies.

Topics:

- Key workloads
- New technologies

Key workloads

The target workloads for the PowerEdge R760xs include virtualization, medium density VM and VDI, and scale-out database. .

New technologies

Table 1. New technologies

Technology	Detailed Description
Up to 2 x 5th Gen Intel(R) Xeon(R) Scalable Processors (Emerald Rapids)	Core count: Up to 28 per processor
	Up to 3.9 GHz
	UPI speed: Up to 3x UPIs/Socket at 12.8 GT/s or 14.4 or 16 GT/s or 20 GT/s
	Maximum number of PCIe lanes: Integrated 80 PCIe 5.0 lanes @ 32 GT/s PCIe Gen5
	Maximum TDP: 250 W
Up to 2 x 4th Gen Intel(R) Xeon(R) Scalable Processors (Sapphire Rapids)	Core count: Up to 32 per processor UPI speed: Up to 3x UPIs/Socket at 12.8 GT/s or 14.4 or 16 GT/s
	Maximum number of PCIe lanes: Integrated 80 PCIe 5.0 lanes @ 32 GT/s PCIe Gen5
	Maximum TDP: 250 W
DDR5 ECC memory up to 5200 MT/s	Max 8 DIMMs per processor and 16 DIMMs per system
	Supports DDR5 ECC RDIMM
GPUs	Max 2 x 75 W SW GPUs (NVIDIA A2)
Flex I/O	LOM: 2x1GbE with BCM5720 LAN controller
	Rear I/O with: • 1x Dedicated iDRAC Ethernet port (1 GbE) • 1 x USB 3.0 • 1 x USB 2.0 • 1 x VGA port



Table 1. New technologies (continued)

Technology	Detailed Description
	Serial Port option
	Optional OCP Mezz 3.0 (supported by x8 PCIe lanes)
	Front I/O with: • 1 x USB 2.0 • 1x iDRAC Direct (Micro-AB USB) port • 1 x VGA port
CPLD 1-wire	Support payload data of Front PERC, Riser, BP and Rear IO to BOSS-N1 and iDRAC
Dedicated PERC	Front Storage module PERC with Front PERC11, PERC12 and PERC 12.2
Software RAID	OS RAID / S160
Power Supplies	60 mm dimension is the new PSU form factor on 16G design 600 W Platinum 100–240 VAC/ 240 VDC 700 W Titanium 200–240 VAC/ 240 VDC 800 W Platinum 100–240 VAC/ 240 VDC 1100 W DC/-48–(-60) V 1100 W Titanium 100–240 VAC/ 240 VDC 1400 W Titanium 100–240 VAC/ 240 VDC 1400 W Platinum 100–240 VAC/ 240 VDC 1400 W Titanium 277 VAC/ 336 VDC 1800 W Titanium 200–240 VAC/ 240 VDC



System features and generational comparison

The following table shows the comparison between the PowerEdge R760xs with the PowerEdge R750xs.

Table 2. Features comparison

Feature	PowerEdge R760xs	PowerEdge R750xs		
Processor	 Up to 2 x 4th Gen Intel(R) Xeon(R) Scalable Processors (Sapphire Rapids) with up to 32 cores Up to 2 x 5th Gen Intel(R) Xeon(R) Scalable Processors (Emerald Rapids) with up to 28 cores 	Maximum two 3 rd Generation Intel [®] Xeon [®] Scalable processors with maximum 32 cores per processor		
Processor Interconnect	Intel Ultra Path Interconnect (UPI) , up to 3 links per CPU	Intel Ultra Path Interconnect (UPI)		
Memory	16 DDR5 DIMM slots	16 DDR4 DIMM slots		
	Supports RDIMM 1.5 TB max	Supports RDIMM 1 TB max		
	Speed maximum 5200 MT/s for 5th generation	Speed maximum 3200 MT/s		
	and 4800 MT/s for 4th generation processors	Supports registered ECC DDR4 DIMMs only		
	Supports registered ECC DDR5 DIMMs only	Apache Pass : No		
	NVDIMM : No	NVDIMM : No		
Storage Drives	 Front bays: 0 drive bay Maximum 8x 3.5-inch SAS/SATA (HDD/ SSD) max 160 TB Maximum 12x 3.5-inch SAS/SATA (HDD/ SSD) max 240 TB Maximum 8x 2.5-inch SAS/SATA/NVMe (HDD/SSD) max 122.88 TB Maximum 16x 2.5-inch SAS/SATA (HDD/ SSD) max 121.6 TB Maximum 16x 2.5-inch (SAS/SATA) + 8x 2.5- inch (NVMe) (HDD/SSD) max 244.48 TB Rear bays: Maximum 2x 2.5 inch SAS/SATA (NVMe) 	 Front bays: 0 drive bay Maximum 8x 3.5-inch SAS/SATA (HDD/SSD) max 128 TB Maximum 12x 3.5-inch SAS/SATA (HDD/SSD) maximum 192 TB Maximum 8x 2.5-inch SAS/SATA/NVMe (HDD/ SSD) maximum 61.44 TB Maximum 16x 2.5-inch SAS/SATA (HDD/SSD) maximum 122.88 TB Maximum 16x 2.5-inch (SAS/SATA) + 8x 2.5- inch (NVMe) (HDD/SSD) maximum 184.32 TB Rear bays: Maximum 2x 2.5-inch SAS (SATA (NVMe) (HDD/S) 		
	Maximum 2x 2.5-inch SAS/SATA/NVMe (HDD/SSD) max 30.72 TB	• Maximum 2x 2.5-inch SAS/SATA/NVMe (HDD/ SSD) maximum 15.36 TB		
Storage Controllers	Internal controllers: H965i, HBA465i (post RTS), H755, H755N, H355, HBA355i	Internal controllers: PERC H345, PERC H355, PERC H745, PERC H755, PERC H755N, HBA355i		
	Internal Boot: Boot Optimized Storage Subsystem (BOSS N1): HWRAID 2x M.2 SSDs and Internal USB	Internal Boot: Internal Dual SD Module or Boot Optimized Storage Subsystem (BOSS S2): HWRAID 2x M.2 SSDs or Internal USB		
	External: HBA355e, H965e and HBA465e (post	External PERC (RAID): PERC H840, HBA355e		
	Software RAID: S160	Software RAID: S150		



Table 2. Features comparison (continued)

Feature	PowerEdge R760xs	PowerEdge R750xs	
PCIe SSD	Front: Maximum 8 x 2.5-inch (NVMe drives)	Maximum 8 x 2.5-inch (NVMe drives)	
	Rear: Up to 2 x 2.5-inch NVMe		
PCle Slots	Up to 6 PCIe slots (2 x Gen5, 4 x Gen4)	Up to 6 PCle slots (5 x Gen4, 1 x Gen3)	
Embedded NIC (LOM)	2x 1GbE LOM	2x 1GbE LOM	
Networking Options (OCP 3.0)	Rear: 1 x OCP 3.0 (x8 PCle lanes)	Maximum 1 OCP 3.0 (x16 PCIe lanes)	
GPU	Nvidia A2 (60 W, LP)	Not supported	
I/O Ports	Front ports 1x Dedicated iDRAC micro-USB 1x USB 2.0 1x VGA 	Front ports 1x Dedicated iDRAC micro-USB 1x USB 2.0 1x VGA 	
	Rear ports: • 1 x Dedicated iDRAC Ethernet port • 1x USB 2.0 • 1x USB 3.0 • 1x Serial (optional) • 1x VGA • 2x Ethernet	Rear ports: • 1 x Dedicated iDRAC Ethernet port • 1x USB 2.0 • 1x USB 3.0 • 1x Serial (optional) • 1x VGA • 2x Ethernet	
	Internal port: • 1x USB 3.0 (optional)	Internal port: • 1x USB 3.0 (optional)	
Rack Height	2U rack server	2U rack server	
Power Supplies	 1800 W Titanium 200–240 VAC/ 240 VDC 1400 W Platinum 100–240 VAC/ 240 VDC 1400 W Titanium 277 VAC/ 336 VDC 1100 W Titanium 100–240 VAC/ 240 VDC 1100 W DC/-48–(-60) V 800 W Platinum 100–240 VAC/ 240 VDC 700 W Titanium 200–240 VAC/ 240 VDC 600 W Platinum 100–240 VAC/ 240 VDC 	 1800 W Platinum 100-240 VAC/ 240 VDC 1400 W Platinum 100-240 VAC/ 240 VDC 1100 W Titanium 100-240 VAC/ 240 VDC 1100 W DC/-48-(-60) V 800 W Platinum 100-240 VAC/ 240 VDC 700 W Titanium 200-240 VAC/240 VDC 600 W Platinum 100-240 VAC/ 240 VDC 	
System Management	 Lifecycle Controller 3.x OpenManage QuickSync 2.0 OpenManage Enterprise Power Manager Digital License Key iDRAC Direct (dedicated micro-USB port) Easy Restore 	 Lifecycle Controller 3.x OpenManage QuickSync 2.0 OpenManage Enterprise Power Manager Digital License Key iDRAC Direct (dedicated micro-USB port) Easy Restore 	
Availability	Hot-plug drives	Hot-plug drives	
	Hot-plug redundant cooling	Hot-plug redundant cooling	
	Hot-plug redundant power supplies	Hot-plug redundant power supplies	
	BOSS-N1	IDSDM	
		BOSS S2	



Chassis views and features

Topics:

Chassis views



Front view of the system



Figure 1. Front view of 16 x 2.5-inch SAS/SATA + 8 x 2.5-inch NVMe drive system



Figure 2. Front view of 16 x 2.5-inch SAS/SATA drive system

ů l	

Figure 3. Front view of 8 x 2.5-inch SAS/SATA or NVMe drive system



Figure 4. Front view of 12 x 3.5-inch SAS/SATA drive system



	顧問題言	1	國際觀		國際關係	C		ju I
110		10		. 6	國政部回		200 m (1) ==	ũ

Figure 5. Front view of 8 x 3.5-inch SAS/SATA drive system



Figure 6. Front view of no backplane configuration (0 drive system)

Left control panel view



Figure 7. Left control panel

Table 3. Left control panel

ltem	Indicator, button, or connector	lcon	Description
1	Status LED indicators	NA	Indicates the status of the system. For more information, see the Status LED indicators section.
2	System health and system ID	ž.	Indicates the system health. For more information, see the System health and system ID indicator codes section.





Figure 8. Left control panel with optional iDRAC Quick Sync 2 indicator

Table 4. Left control panel with optional iDRAC Quick Sync 2 indicator

ltem	Indicator, button, or connector	lcon	Description
1	Status LED indicators	N/A	Indicates the status of the system. For more information, see the Status LED indicators section.
2	System health and system ID indicator	ż	Indicates the system health. For more information, see the System health and system ID indicator codes section.
3	iDRAC Quick Sync 2 wireless indicator (optional)		Indicates if the iDRAC Quick Sync 2 wireless option is activated. The Quick Sync 2 feature allows management of the system using mobile devices. This feature aggregates hardware/ firmware inventory and various system level diagnostic/error information that can be used in troubleshooting the system. You can access system inventory, Dell Lifecycle Controller logs or system logs, system health status, and also configure iDRAC, BIOS, and networking parameters. You can also launch the virtual Keyboard, Video, and Mouse (KVM) viewer and virtual Kernel- based Virtual Machine (KVM), on a supported mobile device. For more information, see the Integrated Dell Remote Access Controller User's Guide at www.dell.com/poweredgemanuals.

(i) NOTE: For more information about the indicator codes, see the System diagnostics and indicator codes section.

Right control panel view



Figure 9. Right control panel



Table 5. Right control panel

ltem	Indicator or button	lcon	Description
1	Power button	Ċ	Indicates if the system is powered on or off. Press the power button to manually power on or off the system. Image: Note: Press the power button to gracefully shut down an ACPI-compliant operating system.
2	USB 2.0 port	4	The USB port is 4-pin, 2.0-compliant. This port enables you to connect USB devices to the system.
3	iDRAC Direct (Micro-AB USB) port	3.E	 The iDRAC Direct (Micro-AB USB) port enables you to access the iDRAC direct Micro-AB USB features. For more information, see the <i>Integrated Dell Remote Access Controller User's Guide</i> at www.dell.com/poweredgemanuals. NOTE: You can configure iDRAC Direct by using a USB to micro USB (type AB) cable, which you can connect to your laptop or tablet. Cable length should not exceed 3 feet (0.91 meters). Performance could be affected by cable quality.
4	VGA port	ici	Enables you to connect a display device to the system.

Rear view of the system

Figure 10. Rear view of the system



Figure 11. Rear view of the system with no riser and one CPU



Figure 12. Rear view of the system with no riser and two CPUs





Figure 13. Rear view of the system with Riser 1c



Figure 14. Rear view of the system with Riser 1d

Chassis views and features 13



Inside the system



Figure 15. Inside the system without rear drive and riser

- 1. Rear mounted front PERC
- 3. Memory module slots
- 5. Intrusion switch
- 7. PSU 1 and PSU 2
- 9. Processor heat sink
- 11. NVMe backplane

- 2. Cooling fan assembly
- 4. Power interposer board
- 6. OCP
- 8. System board
- 10. SAS/SATA backplane
- 12. Information tag




Figure 16. Inside the system with rear drive cage and riser

- 1. Rear mounted front PERC
- 3. Memory module slots
- 5. Intrusion switch
- 7. PSU 1 and PSU 2
- 9. System board
- 11. SAS/SATA backplane
- 13. Information tag

- 2. Cooling fan assembly
- 4. Power interposer board
- 6. Rear drive cage
- 8. Riser
- 10. Processor heat sink
- 12. NVMe backplane

System diagnostics and indicator codes

The diagnostic indicators on the system front panel display system status during system startup.

Status LED indicators

(i) NOTE: The indicators display solid amber if any error occurs.

Chassis views and features 15





Figure 17. Status LED indicators

Table 6.	Status	LED	indicators	and	descri	ptions
----------	--------	-----	------------	-----	--------	--------

lcon	Description	Condition	Corrective action			
0	Drive indicator	The indicator turns solid amber if there is a drive error.	 Check the System Event Log to determine if the drive has an error. Run the appropriate Online Diagnostics test. Restart the system and run embedded diagnostics (ePSA). If the drives are configured in a RAID array, restart the system, and enter the host adapter configuration utility program. 			
8	Temperature indicator	The indicator turns solid amber if the system experiences a thermal error (for example, the ambient temperature is out of range or there is a fan failure).	 Ensure that none of the following conditions exist: A cooling fan has been removed or has failed. System cover, air shrouds, or back filler bracket has been removed. Ambient temperature is too high. External airflow is obstructed. If the problem persists, see the Getting help section. 			
	Electrical indicator	The indicator turns solid amber if the system experiences an electrical error (for example, voltage out of range, or a failed power supply unit (PSU) or voltage regulator).	Check the System Event Log or system messages for the specific issue. If it is due to a problem with the PSU, check the LED on the PSU. Reseat the PSU. If the problem persists, see the Getting help section.			
ø	Memory indicator	The indicator turns solid amber if a memory error occurs.	Check the System Event Log or system messages for the location of the failed memory. Reseat the memory module. If the problem persists, see the Getting help section.			
D	PCIe indicator	The indicator turns solid amber if a PCle card experiences an error.	Restart the system. Update any required drivers for the PCle card. Reinstall the card.If the problem persists, see the Getting help section.If the problem persists, see the Expansion cards and expansion card risers > Expansion card installation guidelines section.			



System health and system ID indicator codes

The system health and system ID indicator is located on the left control panel of the system.



Figure 18. System health and system ID indicator

Table 7. System health and system ID indicator codes

System health and system ID indicator code	Condition
Solid blue	Indicates that the system is powered on, is healthy, and system ID mode is not active. Press the system health and system ID button to switch to system ID mode.
Blinking blue	Indicates that the system ID mode is active. Press the system health and system ID button to switch to system health mode.
Solid amber	Indicates that the system is in fail-safe mode. If the problem persists, see the Getting help section.
Blinking amber	Indicates that the system is experiencing a fault. Check the System Event Log for specific error messages. For information about the event and error messages generated by the system firmware and agents that monitor system components, go to qrl.dell.com > Look Up > Error Code, type the error code, and then click Look it up.



QR code for PowerEdge R760xs system resources



Figure 19. QR code for PowerEdge R760xs system

18 Chassis views and features

Inserido ao protocolo **22.951.206-4** por: **Pedro Henrique Golin Linhares** em: 31/03/2025 12:49. A autenticidade deste documento pode ser validada no endereço: https://www.eprotocolo.pr.gov.br/spiweb/validarDocumento com o código: fc1c44dfe0686d5547ffd691199be699.





Topics:

- Processor features
- Chipset

Processor features

The Intel 4th and 5th Generation Xeon[®] Scalable Processors stack is the next-generation data center processor offering with significant performance increases, integrated acceleration, and next-generation memory and I/O. Sapphire Rapids and Emerald Rapids accelerate customer usage with unique workload optimizations and provide the following feature improvements.

- Faster UPI with up to three Intel Ultra Path Interconnect (Intel UPI) at up to 20 GT/s, increasing multisocket bandwidth.
- More, faster I/O with PCI Express 5 and up to 80 lanes (per CPU)
- Enhanced Memory Performance with DDR5 support and memory speed up to 5200 MT/s in one DIMM per channel (1DPC).
- New onboard accelerators for data analytics, networking, storage, crypto, and data compression
- Enhanced security for virtualized environments with Intel Trust Domain Extensions (IntelR TDX) for confidential computing

Supported processors

The following table shows the Intel Sapphire Rapids and Emerald Rapids SKUs that are supported on the R760xs.

Processor	Clock Speed (GHz)	Cache (M)	UPI (GT/s)	Cores	Turbo	Memory Speed (MT/s)	TDP
6534	3.9	22.5	20	8	Turbo	4800	195 W
6526Y	2.8	37.5	20	16	Turbo	5200	195 W
5512U	2.1	52.5	N/A	28	Turbo	4800	185 W
4514Y	2.0	30	16	12	Turbo	4400	150 W
4510	2.4	30	16	12	Turbo	4400	150 W
4509Y	2.6	23	16	8	Turbo	4400	125 W
6448Y	2 .2	60	16	32	Turbo	4800	225 W
6442Y	2.6	45	16	24	Turbo	4800	225 W
6438Y+	2	60	16	32	Turbo	4800	205 W
6426Y	2.6	30	16	16	Turbo	4800	185 W
6414U	2	60	16	32	Turbo	4800	250 W
5420+	2	53	16	28	Turbo	4400	205 W
5418Y	2	45	16	24	Turbo	4400	185 W
5416S	2	30	16	16	Turbo	4400	150 W
5415+	2.9	15	16	8	Turbo	4400	150 W
5412U	2.1	45	16	24	Turbo	4400	185 W
4416+	2	38	16	20	Turbo	4000	165 W

Table 8. Supported Processors for R760xs

Processor 19



Table 8. Supported Processors for R760xs (continued)

Processor	Clock Speed (GHz)	Cache (M)	UPI (GT/s)	Cores	Turbo	Memory Speed (MT/s)	TDP
4410Y	2	23	16	12	Turbo	4000	150 W
4410T	2.7	27	16	10	Turbo	4000	150 W
3408U	1.8	15	16	8	Turbo	4000	125 W

Chipset

The system supports Intel[®] C741 series chipset. DMI - 3.0 speed (port width x8, x4) USB ports - up to 10 superspeed (USB 3.1), 14 highspeed (USB 2.0)

SATA ports - up to 20 SATA port

PCIe Express - Up to 20 lanes, PCIe 3.0

Chipset features

- PCI-E interfaces
 - \circ $\;$ Integrated PCI Express Gen5 for improved bandwidth and connectivity
 - \circ $\,$ Up to 80 lanes per processor
 - Connect PCIe x1 to iDRAC- integrated VGA chip
- Integrated USB maximum of 10 SuperSpeed (USB 3.1), 14 highspeed (USB 2.0)
 - \circ $\,$ One front port (USB 2.0 / Right front I/O) $\,$
 - Two rear ports (USB 2.0/3.0)



Memory subsystem

Topics:

- Supported memory
- General memory module installation guidelines

Supported memory

Table 9. Memory technology comparison

Feature	PowerEdge R760xs (DDR5)
DIMM type	RDIMM
Transfer speed	5200 MT/s (1DPC) () NOTE: Maximum DIMM transfer speed support depends on CPU SKU and DIMM population.
Voltage	1.1 V

Table 10. Supported memory matrix

DIMM type	Rank	Capacity	DIMM rated voltage	Operating Speed	
			and speed	1 DIMM per channel (DPC)	
RDIMM	1 R	16 GB	DDR5 (1.1 V), 4800 MT/s	Up to 4800 MT/s	
			DDR5 (1.1 V), 5600 MT/s	op to 3200 Mit/s	
	2 R	32 GB, 64 GB, 96 GB	DDR5 (1.1 V), 4800 MT/s DDR5 (1.1 V), 5600 MT/s	Up to 4800 MT/s Up to 5200 MT/s	

(i) NOTE: 5600 MT/s RDIMMs are applicable for 5th Gen IntelR XeonR Scalable Processors.

(i) NOTE: The processor may reduce the performance of the rated DIMM speed.

General memory module installation guidelines

To ensure optimal performance of your system, observe the following general guidelines when configuring your system memory. If your system's memory configuration fails to observe these guidelines, your system might not boot, stop responding during memory configuration, or operate with reduced memory.

The memory bus may operate at speeds of 5200 MT/s, 4800 MT/s, 4400 MT/s or 4000 MT/s depending on the following factors:

- System profile selected (for example, Performance, Performance Per Watt Optimized (OS), or Custom [can be run at high speed or lower])
- Maximum supported DIMM speed of the processors



• Maximum supported speed of the DIMMs

(i) NOTE: MT/s indicates DIMM speed in MegaTransfers per second.

The system supports Flexible Memory Configuration, enabling the system to be configured and run in any valid chipset architectural configuration. The following are the recommended guidelines for installing memory modules:

- All DIMMs must be DDR5.
- Memory mixing is not supported for different DIMM capacities.
- If memory modules with different speeds are installed, they operate at the speed of the slowest installed memory module(s).
- Populate memory module sockets only if a processor is installed.
 - For single-processor systems, sockets A1 to A8 are available.
 - $\circ~$ For dual-processor systems, sockets A1 to A8 and sockets B1 to B8 are available.
 - $\circ~$ A minimum of 1 DIMM must be populated for each installed processor.
- In **Optimizer Mode**, the DRAM controllers operate independently in the 64-bit mode and provide optimized memory performance.

Table 11. Memory population rules

Processor	Configuration	Memory population	Memory population information
Single processor	Optimizer (Independent channel) population order	A{1}, A{2}, A{3}, A{4}, A{5}, A{6}, A{7}, A{8}	1, 2, 4, 6, 8 DIMMs are allowed.
Dual processor (Start with processor1. Processor 1 and processor 2 population should match)	Optimizer (Independent channel) population order	A{1}, B{1}, A{2}, B{2}, A{3}, B{3}, A{4}, B{4}, A{5}, B{5}, A{6}, B{6}, A{7}, B{7} A{8}, B{8}	 2, 4, 8, 12, 16 DIMMs are supported per system . i) NOTE: Optimizer population order is not traditional for 8 and 16 DIMMs installations for dual processor.

• Always populate memory channels identically with equal DIMMs for best performance.

• Supported RDIMM configurations are 1, 2, 4, 6, and 8 DIMMs per processor.

• Supported 96 GB RDIMM configurations are 1, 6 and 8 DIMMs per processor.

• Populate eight equal memory modules per processor (one DIMM per channel) at a time to maximize performance.

NOTE: Equal memory modules refer to DIMMs with identical electrical specification and capacity that may be from different vendors.





Topics:

- Storage controllers
- Supported Drives
- External Storage

Storage controllers

Dell RAID controller options offer performance improvements, including the fPERC solution. fPERC provides a base RAID HW controller without consuming a PCIe slot by using a small form factor and high-density connector to the base planar.

16G PERC Controller offerings are a heavy leverage of 15G PERC family. The Value and Value Performance levels carry over to 16G from 15G. New to 16G is the Avenger-based Premium Performance tier offering. This high-end offering drives IOPs performance and enhanced SSD performance.

Table 12. PERC Series controller offerings

Performance Level	Controller and Description
Entry	S160
Value	H355, HBA355e, HBA355i, (internal/external)
Value Performance	H755, H755N
Premium Performance	Н965і,
	Avenger 1
	Memory: 8GB DDR4 NV cache
	72-bit memory 2133 MHz
	Low profile form factors
	Dual A15 1.2 GHz CPU
	X8PCle 3.0, x8 12Gb SAS

() NOTE: For more information about the features of the Dell PowerEdge RAID controllers (PERC), Software RAID controllers, or BOSS card, and on deploying the cards, see the storage controller documentation at www.dell.com/storagecontrollermanuals.

(i) NOTE: From December 2021, H355 replaces H345 as the entry raid controller. H345 is deprecated in January 2022.

Storage controller feature matrix

Table 13. Storage controller feature matrix

Model & Form Factors	Interface Support	PCI Suppo rt	SAS Connection	Cach e Mem ory Size	Write Back Cache	RAID Levels	Max Drive Support	RAID Support
PowerEdge Server-Storage Controllers (PERC) Series 12								



Table 13. Storage controller feature matrix (continued)

Model & Form Factors	Interface Support	PCI Suppo rt	SAS Connection	Cach e Mem ory Size	Write Back Cache	RAID Levels	Max Drive Support	RAID Support
H965i Front	24Gb/s SAS 6Gb/s SAS/SATA Gen3 (8 GT/s) NVMe Gen4 (16 GT/s) NVMe	PCle Gen 4	16 ports/lanes - 2x8 Internal	8GB NV	Flash Backed Cache	0,1,5,6,10,50 ,60	16	Hardware
H965i Adapter	24Gb/s SAS 6Gb/s SAS/SATA Gen3 (8 GT/s) NVMe Gen4 (16 GT/s) NVMe	PCle Gen 4	16 ports/lanes - 2x8 Internal	8GB NV	Flash Backed Cache	0,1,5,6,10,50 ,60	16	Hardware
S160 Software RAID	Gen4 (16 GT/s) NVMe	PCle Gen 4	N/A	No Cach e	No Cache	0,1,5,10	8	Software RAID - Windows only
	PowerE	dge Serv	er-Storage Cont	rollers (PERC & SAS H	HBA) Series 11		
H755 Adapter	12Gb/s SAS 6Gb/s SAS/SATA 3Gb/s SAS/SATA Gen3 (8 GT/s) NVMe Gen4 (16 GT/s) NVMe	PCle Gen 4	16 ports- 2x8 Internal	8GB NV	Flash Backed Cache	0,1,5,6,10,50	16/ controller 50 with SAS Expander	Hardware
H755N Front (NVMe Only)	Gen3 (8 GT/s) NVMe Gen4 (16 GT/s) NVMe	PCle Gen 4	16 ports- 2x8 Internal	8GB NV	Flash Backed Cache	0,1,5,6,10,50 ,60	8/ controller	Hardware
H755 Front (SAS/ SATA only)	12Gb/s SAS 6Gb/s SAS/SATA 3Gb/s SAS/SATA	PCle Gen 4	16 ports- 2x8 Internal	8GB NV	Flash Backed Cache	0,1,5,6,10,50 ,60	16/ controller 50 with SAS Expander	Hardware
HBA355i Adapter	12Gb/s SAS 6Gb/s SAS/SATA 3Gb/s SAS/SATA	PCle Gen 4	16 ports- 2x8 Internal	N/A	N/A	N/A	16/ controller 50 with SAS Expander	N/A
HBA355i Front	12Gb/s SAS 6Gb/s SAS/SATA 3Gb/s SAS/SATA	PCle Gen 4	16 ports- 2x8 Internal	N/A	N/A	N/A	16/ controller 50 with SAS Expander	N/A

24 Storage

Inserido ao protocolo **22.951.206-4** por: **Pedro Henrique Golin Linhares** em: 31/03/2025 12:49. A autenticidade deste documento pode ser validada no endereço: https://www.eprotocolo.pr.gov.br/spiweb/validarDocumento com o código: fc1c44dfe0686d5547ffd691199be699.



Table 13. Storage controller feature matrix (continued)

Model & Form Factors	Interface Support	PCI Suppo rt	SAS Connection	Cach e Mem ory Size	Write Back Cache	RAID Levels	Max Drive Support	RAID Support
HBA355e Adapter	12Gb/s SAS 6Gb/s SAS/SATA 3Gb/s SAS/SATA	PCle Gen 4	16 ports- 4x4 external	NZA	N/A	N/A	240	N/A
H355 Adapter	12Gb/s SAS 6Gb/s SAS/SATA	PCle Gen 4	16 ports- 2x8 Internal	No Cach e	No Cache	0,1, 10	Up to 32 RAID, or 32 Non- RAID	Hardware
H355 Front	12Gb/s SAS 6Gb/s SAS/SATA	PCle Gen 4	16 ports- 2x8 Internal	No Cach e	No Cache	0,1, 10	Up to 32 RAID, or 32 Non- RAID	Hardware

() NOTE:

- 1. RAID 5/50 removed from entry RAID card
- **2.** SWRAID support for Linus provides a pre-boot configuration utility to configure MDRAID and degraded boot capability.
- **3.** For information, post-RTS, see the Storage controller documentation at www.dell.com/stroagecontrollermanuals.

This document is updated as changes happen, so for the latest version be sure to bookmark it rather than downloading an offline copy or refer to the Storage Controller Matrix on sales portal.

Server storage controllers User Guide

• Server-Storage Controllers User's Guides, click here

RAID - Redundant Array of Independent Disks

• Link to Help Me Choose: RAID Configuration here

Datasheets and PERC performance scaling decks

- Resource Page for Server-Storage (Sales Portal) click here
- PERC & SAS HBA Datasheets (To be updated)

Boot Optimized Storage Solution (BOSS)

BOSS is a RAID solution that is designed to boot operating systems and segregate operating system boot drives from data on server-internal storage.



BOSS feature matrix

Table 14. BOSS feature matrix

BOSS card	Drive Size	RAID levels	Stripe size	Virtual disk cache functio n	Maxim um numbe r of virtual disks	Maxim um numbe r of drives suppor ted	Drive types	PCIe suppor t	Disk cache policy	Suppor t for Non- RAID disks	Crypto graphi c digital signatu re to verify firmwa re payloa d	Hot Plug
BOSS- N1 Monolit hic	M.2 devices are read- intensiv e with 480 GB or 960 GB capacit y	RAID1 and RAID0	Support s default 64K stripe size only	None	1	2	M.2 NVMe SSDs	Gen3	Drive default	No	Yes	Yes

BOSS-N1

BOSS-N1 is offered as a means of booting 16G servers to a full OS when the target OS is a full OS (not just a hypervisor), or the user does not wish to trade off standard hot plug drive slots for OS install.

The HW RAID BOSS-N1 card is a RAID controller with a limited feature set that presents M.2 NVMe-only SSDs as either a RAID0 disk or a single RAID1 volume with 2 disks. BOSS-N1 enables support for 480/960 GB Disks from Factory Install.

Hardware: BOSS-N1 Controller and Carrier (x2)

Reliability: Enterprise-Class M.2 NVMe SSDs

Supports dual 80 mm, Read Intensive (1DWPD), M.2 devices 480 GB/960 GB Standard - 1.92 TB QNS

Accessibility: Front Facing

Serviceability: Full Hot-Plug Support

Supports Hardware RAID1 and RAID0

Supports UEFI boot

Marvell 88NR2241 NVMe RAID Controller

Controlled Firmware Upgrade through iDRAC



Figure 20. BOSS-N1 Controller



Datasheets

• BOSS-N1 (to be updated)

BOSS User Guides

• BOSS-N1

Supported Drives

The table shown below lists the internal drives supported by the R760xs. Refer to Agile for the latest SDL

Table 15. Supported Drives

Form Factor	Туре	Speed	Rotational Speed	Capacities
2.5 inches	vSAS	12 Gb	SSD	1.92 TB, 3.84 TB, 960 GB, 7.62 TB
2.5 inches	SAS	24 Gb	SSD	1.92 TB, 1.6 TB, 800 GB, 3.84 TB, 960 GB, 7.68 TB
2.5 inches	SATA	6 Gb	SSD	1.92 TB, 480 GB, 960 GB, 3.84 TB
2.5 inches	NVMe	Gen4	SSD	1.6 TB, 3.2 TB, 6.4 TB, 1.92 TB, 3.84 TB, 15.63 TB, 7.68 TB
2.5 inches	DC NVMe	Gen4	SSD	3.84 TB, 960 GB
2.5 inches	SAS	12 Gb	10 K	600 GB, 1.2 TB, 2.4 TB
3.5 inches	SATA	6 Gb	7.2 K	2 TB, 4 TB, 8 TB, 12 TB, 16 TB, 20 TB
3.5 inches	SAS	12 Gb	7.2 K	2 TB, 4 TB, 8 TB, 12 TB, 16 TB, 20 TB

External Storage

The R760xs support the external storage device types listed in the table below.

Table 16. Support External Storage Devices

Device Type	Description	
External Tape	Supports connection to external USB tape products	
NAS/IDM appliance software	Supports NAS software stack	
JBOD	Supports connection to 12 Gb MD-series JBODs	





Topics:

- Overview
- OCP 3.0 support

Overview

PowerEdge offers a wide variety of options to get information moving to and from our servers. Industry best technologies are chosen, and systems management features are added by our partners to firmware to tie in with iDRAC. These adapters are rigorously validated for worry-free, fully supported use in Dell servers.

OCP 3.0 support

Table 17. OCP 3.0 feature list

Feature	OCP 3.0
Form factor	SFF
PCle Gen	Gen4
Max PCle width	x8
Max no. of ports	4
Port type	BT/SFP/SFP+/SFP28
Max port speed	25 GbE
NC-SI	Yes
SNAPI	No
WoL	Yes
Power consumption	15 W–35 W

Supported OCP cards

Table 18. Supported OCP cards

Form factor	Vendor	Port speed	Port type	Port count	DPN
OCP 3.0	Broadcom	10 GbE	BT	2	RN1M5
		25 GbE	SFP28	2	24FG6
		25 GbE	SFP28	4	3Y64D
		1 GbE	BT	4	VJWVJ
		10 GbE	ВТ	4	W5HC8
	Intel	1 GbE	ВТ	4	HY4CV

28 Networking



Table 18. Supported OCP cards (continued)

Form factor	Vendor	Port speed	Port type	Port count	DPN
		10 GbE	BT	2	F6X1R
		10 GbE	BT	4	XC0M4
		25 GbE	SFP28	2	PWH3C
		25 GbE	SFP28	4	Y4VV5

OCP NIC 3.0 vs. rack Network Daughter Card comparisons

Table 19. OCP 3.0, 2.0, and rNDC NIC comparison

Form Factor	Dell rNDC	OCP 2.0 (LOM Mezz)	OCP 3.0	Notes
PCle Gen	Gen 3	Gen 3	Gen 4	Supported OCP3 are SFF (small form factor)
Max PCle Lanes	x8	Up to x16	Up to x8	See server slot priority matrix
Shared LOM	Yes	Yes	Yes	This is iDRAC port redirect
Aux Power	Yes	Yes	Yes	Used for Shared LOM

OCP form factors



Figure 21. OCP 3.0 Small Card Form Factor (LS)

The process of installing the OCP card in R760xs system:

- 1. Open the blue latch on the system board.
- 2. Slide the OCP card into the slot in the system.
- **3.** Push until the OCP card is connected to the connector on the system board.
- **4.** Close the latch to lock the OCP card to the system.





Figure 22. Installing the OCP Card in R760xs

The process of removing the OCP card in R760xs system:

- 1. Open the blue latch to unlock the OCP card.
- 2. Push the OCP card towards the rear end of the system to disconnect from the connector on the system board.
- $\ensuremath{\textbf{3.}}$ Slide the OCP card out of the slot on the system.



Figure 23. Removing the OCP Card in R760xs



Slot priority matrix

For add-in cards that can be mapped to the R760xs and guidelines for installing expansion cards, see the R760xs slot priority matrix file on Sales Portal.

Link:https://www.delltechnologies.com/resources/en-us/auth/products/servers/category.htm

Topics:

• Expansion card installation guidelines

Expansion card installation guidelines



Figure 24. Expansion card slot connectors

- 1. SL9_CPU2_PB5 (PCIe cable connector for Riser 1C and Riser 1D)
- 2. SL10_CPU2_PA5 (PCIe cable connector for Riser 1C and Riser 1D)
- 3. SIG_PWR_0 (Power connector for Riser 1C and Riser 1D)
- **4.** SL11_CPU1_PA6 (PCIe cable connector for Riser 1D)
- 5. SL12_CPU1_PB6 (PCIe cable connector for Riser 1D)

The following table describes the expansion card riser configurations:



Table 20. Expansion card riser configurations

Configuratio ns	Expansion card risers	PCIe Slots	Controlling processor	Height	Length	Slot width	Power
Config 0-1.	No riser	1, 2	Processor 1	Low profile	Half length	x16, x8	75 W
Config 0-2.	No riser	1, 2	Processor 1	Low profile	Half length	x16, x8	75 W
		5, 6	Processor 2	Low profile	Half length	x16, x16	75 W
Config 1.	R1C	1, 2, 3	Processor 1	Low profile	Half length	x16, x8, x16	75 W
		4, 5, 6	Processor 2	Low profile	Half length	x16, x16, x16	75 W
Config 2.	R1D	1, 2, 3, 4	Processor 1	Low profile	Half length	x16, x8, x8, x8	75 W

(i) NOTE: Only one cable riser can be installed at a time in any given configuration.

(i) NOTE: The slots 1, 2, 5 and 6 are Gen4 slots, slot 3 and 4 located on risers are Gen5 slots.



Figure 25. Riser 1C

- 1. Slot 3
- 2. Slot 4

32 Slot priority matrix





Figure 26. Riser 1D

- 1. Slot 3
- 2. Slot 4

(i) NOTE: The expansion-card slots are not hot-swappable.

The following table provides guidelines for installing expansion cards to ensure proper cooling and mechanical fit. The expansion cards with the highest priority should be installed first using the slot priority indicated. All the other expansion cards should be installed in the card priority and slot priority order.

Table 21. Configuration 0-1: No riser configuration

Card type	Slot priority	Maximum number of cards
Dell Serial port module (LP)	2	1
fPERC	Integrated slot	1
InterN/AI PERC adapter	1	1
Dell ExterN/Al Adapter	2, 1	2
Mellanox (NIC: 400Gb)	Not supported	N/A
Mellanox (NIC: 200Gb)	Not supported	N/A
Mellanox (NIC: 100Gb)	1	1
Mellanox HDR100 VPI	1	1
Mellanox HDR VPI	1	1
Broadcom (NIC: 100Gb)	1	1
Intel (NIC: 100Gb)	1	1
Broadcom (SFP: 25Gb)	2, 1	2
Intel (NIC: 25Gb)	2, 1	2
Qlogic (NIC: 25Gb)	Not supported	N/A
Qlogic (NIC: 10Gb)	Not supported	N/A
SolarFlare (NIC: 25Gb)	Not supported	N/A
Broadcom (HBA: FC64)	2,1	2
Broadcom (HBA: FC32)	2, 1	2
Marvell (HBA: FC32)	2, 1	2



Table 21. Configuration 0-1: No riser configuration (continued)

Card type	Slot priority	Maximum number of cards
Emulex (HBA: FC32)	Not supported	N/A
Avago (HBA: FC16)	Not supported	N/A
Qlogic (HBA: FC16)	Not supported	N/A
Broadcom (NIC: 10Gb)	2, 1	2
Intel (NIC: 10Gb)	2, 1	2
Qlogic (NIC: 10Gb)	Not supported	N/A
Broadcom (NIC: 1Gb)	2, 1	N/A
Intel (NIC: 1Gb)	2, 1	2
Intel (OCP: 100Gb)	Integrated slot	1
Broadcom (OCP: 25Gb)	Integrated slot	1
Intel (OCP: 25Gb)	Integrated slot	1
Marvell (OCP: 25Gb)	Not supported	N/A
SolarFlare (OCP: 25Gb)	Not supported	N/A
Broadcom (OCP: 10Gb)	Integrated slot	1
Marvell (OCP: 10Gb)	Not supported	N/A
Intel (OCP: 10Gb)	Not supported	N/A
Broadcom (OCP: 1Gb)	Integrated slot	1
Intel (OCP: 1Gb)	Not supported	N/A
Dell BOSS N1 Module	Integrated slot	1

Table 22. Configuration 0-2: No riser configuration

Card type	Slot priority	Maximum number of cards
Dell Serial port module (LP)	2	1
fPERC	Integrated slot	1
InterN/AI PERC adapter	1	1
Dell ExterN/Al Adapter	6, 2, 1, 5	4
Mellanox (NIC: 400Gb)	Not supported	N/A
Mellanox (NIC: 200Gb)	Not supported	N/A
Mellanox (NIC: 100Gb)	6, 1, 5	3
Mellanox HDR100 VPI	6, 1, 5	3
Mellanox HDR VPI	6, 1, 5	3
Broadcom (NIC: 100Gb)	6, 1, 5	3
Intel (NIC: 100Gb)	6, 1, 5	3
Broadcom (SFP: 25Gb)	6, 2, 1, 5	4
Intel (NIC: 25Gb)	2, 1	2
Qlogic (NIC: 25Gb)	Not supported	N/A
Qlogic (NIC: 10Gb)	Not supported	N/A
SolarFlare (NIC: 25Gb)	Not supported	N/A

34 Slot priority matrix

Inserido ao protocolo **22.951.206-4** por: **Pedro Henrique Golin Linhares** em: 31/03/2025 12:49. A autenticidade deste documento pode ser validada no endereço: https://www.eprotocolo.pr.gov.br/spiweb/validarDocumento com o código: fc1c44dfe0686d5547ffd691199be699.



Table 22. Configuration 0-2: No riser configuration (continued)

Card type	Slot priority	Maximum number of cards
Broadcom (HBA: FC64)	6, 2, 1, 5	4
Broadcom (HBA: FC32)	6, 2, 1, 5	4
Marvell (HBA: FC32)	6, 2, 1, 5	4
Emulex (HBA: FC32)	Not supported	N/A
Avago (HBA: FC16)	Not supported	N/A
Qlogic (HBA: FC16)	Not supported	N/A
Broadcom (NIC: 10Gb)	6, 2, 1, 5	4
Intel (NIC: 10Gb)	6, 2, 1, 5	4
Qlogic (NIC: 10Gb)	Not supported	N/A
Broadcom (NIC: 1Gb)	6, 2, 1, 5	4
Intel (NIC: 1Gb)	6, 2, 1, 5	4
Intel (OCP: 100Gb)	Integrated slot	1
Broadcom (OCP: 25Gb)	Integrated slot	1
Intel (OCP: 25Gb)	Integrated slot	1
Marvell (OCP: 25Gb)	Not supported	N/A
SolarFlare (OCP: 25Gb)	Not supported	N/A
Broadcom (OCP: 10Gb)	Integrated slot	1
Marvell (OCP: 10Gb)	Not supported	N/A
Intel (OCP: 10Gb)	Not supported	N/A
Broadcom (OCP: 1Gb)	Integrated slot	1
Intel (OCP: 1Gb)	Not supported	N/A
Dell BOSS N1 Module	Integrated slot	1

Table 23. Configuration 1: R1C

Card type	Slot priority	Maximum number of cards	
Dell Serial port module (LP)	2	1	
InterN/AI PERC adapter	1	1	
Dell exterN/AI PERC adapter	6, 2, 1, 3, 5, 4	6	
12Gbps SAS HBA	1	1	
Mellanox (NIC: 400Gb)	4, 3	2	
Mellanox (NIC: 200Gb)	4, 3	2	
Broadcom (NIC: 100Gb)	6, 1, 3, 5, 4	5	
Intel (NIC: 100Gb)	6, 1, 3, 5, 4	5	
Mellanox (NIC: 100Gb)	6, 1, 3, 5, 4	5	
Mellanox HDR100 VPI	6, 1, 3, 5, 4	5	
Mellanox HDR VPI	6, 1, 3, 5, 4	5	
Intel (NIC: 25Gb)	6, 2, 1, 3, 5, 4 6		
Mellanox (NIC: 25Gb)	6, 1, 3, 5, 4	5	



Table 23. Configuration 1: R1C (continued)

Card type	Slot priority	Maximum number of cards	
Qlogic (NIC: 25Gb)	Not Supported	N/A	
Broadcom (NIC: 10Gb)	6, 2, 1, 3, 5, 4	6	
Broadcom (NIC: 25Gb)	6, 2, 1, 3, 5, 4	6	
SolarFlare (NIC: 25Gb)	Not Supported	N/A	
Broadcom (HBA: FC64)	6, 2, 1, 3, 5, 4	6	
Broadcom (HBA: FC32)	6, 2, 1, 3, 5, 4	6	
QLogic (HBA: FC32)	6, 2, 1, 3, 5, 4	6	
Marvell (HBA: FC32)	6, 2, 1, 3, 5, 4	6	
Emulex (HBA: FC32)	Not Supported	N/A	
Avago (HBA: FC16)	Not Supported	N/A	
QLogic (HBA: FC16)	Not Supported	N/A	
Intel (NIC: 10Gb)	6, 2, 1, 3, 5, 4	6	
Qlogic (NIC: 10Gb)	Not Supported	N/A	
Intel (NIC: 1Gb)	6, 2, 1, 3, 5, 4	6	
Intel (OCP: 100Gb)	Integrated slot	1	
Broadcom (OCP: 25Gb)	Integrated slot	1	
Intel (OCP: 25Gb)	Integrated slot	1	
Marvell (OCP: 25Gb)	Not supported	N/A	
SolarFlare (OCP: 25Gb)	Not supported	N/A	
Broadcom (OCP: 10Gb)	Integrated slot	1	
Marvell (OCP: 10Gb)	Not supported	N/A	
Intel (OCP: 10Gb)	Not supported	N/A	
Broadcom (OCP: 1Gb)	Integrated slot	1	
Intel (OCP: 1Gb)	Not supported	N/A	
Dell BOSS N1 Module	Integrated slot	1	

Table 24. Configuration 2: R1D

Card type	Slot priority	Maximum number of cards	
Dell Serial port module (LP)	2	1	
InterN/AI PERC adapter	1	1	
Dell exterN/AI PERC adapter	4, 3, 2, 1	4	
12Gbps SAS HBA	1	1	
Mellanox (NIC: 400Gb)	Not supported	N/A	
Mellanox (NIC: 200Gb)	Not supported	N/A	
Broadcom (NIC: 100Gb)	1	1	
Intel (NIC: 100Gb)	1	1	
Mellanox (NIC: 100Gb)	1	1	
Mellanox HDR100 VPI	1	1	



Table 24. Configuration 2: R1D (continued)

Card type	Slot priority	Maximum number of cards	
Mellanox HDR VPI	1	1	
Intel (NIC: 25Gb)	4, 3, 2, 1	4	
Mellanox (NIC: 25Gb)	4, 3, 2, 1	4	
Qlogic (NIC: 25Gb)	Not Supported	N/A	
Broadcom (NIC: 10Gb)	4, 3, 2, 1	4	
Broadcom (NIC: 25Gb)	4, 3, 2, 1	4	
SolarFlare (NIC: 25Gb)	Not Supported	N/A	
Broadcom (HBA: FC64)	4, 3, 2, 1	4	
Broadcom (HBA: FC32)	4, 3, 2, 1	4	
QLogic (HBA: FC32)	4, 3, 2, 1	4	
Marvell (HBA: FC32)	4, 3, 2, 1	4	
Emulex (HBA: FC32)	Not Supported	N/A	
Avago (HBA: FC16)	Not Supported	N/A	
QLogic (HBA: FC16)	Not Supported	N/A	
Intel (NIC: 10Gb)	4, 3, 2, 1	4	
Qlogic (NIC: 10Gb)	Not Supported	N/A	
Intel (NIC: 1Gb)	4, 3, 2, 1	4	
Intel (OCP: 100Gb)	Integrated slot	1	
Broadcom (OCP: 25Gb)	Integrated slot	1	
Intel (OCP: 25Gb)	Integrated slot	1	
Marvell (OCP: 25Gb)	Not supported	N/A	
SolarFlare (OCP: 25Gb)	Not supported	N/A	
Broadcom (OCP: 10Gb)	Integrated slot	1	
Marvell (OCP: 10Gb)	Not supported	N/A	
Intel (OCP: 10Gb)	Not supported	N/A	
Broadcom (OCP: 1Gb)	Integrated slot	1	
Intel (OCP: 1Gb)	Not supported	N/A	
Dell BOSS N1 Module	Integrated slot	1	

(i) NOTE: The serial COM card is not a real PCIe add-in card and has a dedicated slot on the system board.



Power, thermal, and acoustics

PowerEdge servers have an extensive collection of sensors that automatically track thermal activity, which helps to regulate temperature by reducing server noise and power consumption. The table below lists the tools and technologies Dell offers to lower power consumption and increase energy efficiency.

Topics:

- Power
- Thermal
- Acoustics

Power

Table 25. Power tools and technologies

Feature	Description
Power Supply Units(PSU) portfolio	Dell's PSU portfolio includes intelligent features such as dynamically optimizing efficiency while maintaining availability and redundancy. Find additional information in the Power supply units section.
Tools for right sizing	Enterprise Infrastructure Planning Tool (EIPT) is a tool that can help you determine the most efficient configuration possible. With Dell's EIPT, you can calculate the power consumption of your hardware, power infrastructure, and storage at a given workload. Learn more at www.dell.com/calc.
Industry Compliance	Dell's servers are compliant with all relevant industry certifications and guide lines, including 80 PLUS, Climate Savers and ENERGY STAR.
Power monitoring accuracy	PSU power monitoring improvements include:
	 Dell's power monitoring accuracy is currently 1%, whereas the industry standard is 5% More accurate reporting of power Better performance under a power cap
Power capping	Use Dell's systems management to set the power cap limit for your systems to limit the output of a PSU and reduce system power consumption. Dell is the first hardware vendor to leverage Intel Node Manager for circuit-breaker fast capping.
Systems Management	iDRAC Enterprise and Datacenter provides server-level management that monitors, reports and controls power consumption at the processor, memory and system level.
	Dell OpenManage Power Center delivers group power management at the rack, row, and data center level for servers, power distribution units, and uninterruptible power supplies.
Active power management	Intel Node Manager is an embedded technology that provides individual server-level power reporting and power limiting functionality. Dell offers a complete power management solution comprised of Intel Node Manager accessed through Dell iDRAC9 Datacenter and OpenManage Power Center that allows policy-based management of power and thermal at the individual server, rack, and data center level. Hot spare reduces power consumption of redundant power supplies. Thermal control off a speed optimizes the thermal settings for your environment to reduce fan consumption and lower system power consumption. Idle power enables Dell servers to run as efficiently when idle as when at full workload.
Rack infrastructure	Dell offers some of the industry's highest-efficiency power infrastructure solutions, including:



Table 25. Power tools and technologies (continued)

Feature	Description
	 Power distribution units (PDUs) Uninterruptible power supplies (UPSs) Energy Smart containment rack enclosures Find additional information at: https://www.delltechnologies.com/en-us/servers/power-and-cooling.htm.

PSU specifications

The PowerEdge R760xs system supports up to two AC or DC power supply units (PSUs).

Table 26. R760xs PSU specifications	Table	26.	R760xs	PSU s	specifications	
-------------------------------------	-------	-----	---------------	-------	----------------	--

PSU Class Heat dissipation		Frequ AC Voltage ency		e	DC -48 V DC Voltage		277 V AC and HVDC		
		(maximum) (BTU/ hr)	(HZ)	Low Line AC (100– 120 V)	High Line AC (200– 240 V)	240 V DC	-40 V to -72 V DC	277 V AC (249 V AC– 305 V AC)	336 V (260 V DC-400 V DC)
600 W mixed mode	Platinu m	2250	50/60	600 W	600 W	600 W	N/A	N/A	N/A
700 W mixed mode HLAC	Titaniu m	2625	50/60	N/A	700 W	700 W	N/A	N/A	N/A
800 W mixed mode	Platinu m	3000	50/60	800 W	800 W	800 W	N/A	N/A	N/A
1100 W -48 V DC	N/A	4265	N/A	N/A	N/A	N/A	1100 W	N/A	N/A
1100 W mixed mode	Titaniu m	4125	50/60	1050 W	1100 W	1100 W	N/A	N/A	N/A
1400 W mixed mode	Titaniu m	5250	50/60	1050 W	1400 W	1400 W	N/A	N/A	N/A
1400 W mixed mode	Platinu m	5250	50/60	1050 W	1400 W	1400 W	N/A	N/A	N/A
1400 W 277 V AC and HVDC	Titaniu m	5250	50/60	N/A	N/A	N/A	N/A	1400 W	1400 W
1800 W mixed mode HLAC	Titaniu m	6610	50/60	N/A	1800 W	1800 W	N/A	N/A	N/A

(i) NOTE: Heat dissipation is calculated using the PSU wattage rating.

(i) NOTE: HLAC stands for High-Line AC, with a range of 200 - 240V AC. HVDC stands for High-Voltage DC, with 336 V DC.

NOTE: When selecting or upgrading the system configuration, to ensure optimum power utilization, verify the system power consumption with the Enterprise Infrastructure Planning Tool available at Dell.com/calc.

NOTE: If a system with AC 1400 W or 1100 W PSUs operates at low line 100-120 Vac, and then the power rating per PSU is degraded to 1050 W.





Figure 27. PSU power cords



Figure 28. APP 2006G1 power cord



Figure 29. Lotes DC PSU connector

Table 27. PSU power cords

Form factor	Output	Power cord
Redundant 60 mm	600 W Mixed Mode	C13
	700 W Mixed Mode HLAC	C13
	800 W Mixed Mode	C13
	1100 W Mixed Mode	C13
	1100 W -48 V DC	Lotes DC PSU connector
	1400 W Mixed Mode	C13
	1400 W 277 VAC and 336 VDC	APP 2006G1
	1800 W Mixed Mode HLAC	C15

(i) NOTE: C13 power cord combined with C14 to C15 jumper power cord can be used to adapt 1800 W PSU.



Thermal

PowerEdge servers have an extensive collection of sensors that automatically track thermal activity, which helps regulate temperature thereby reducing server noise and power consumption.

Acoustics

Acoustical configurations of R760xs

Dell PowerEdge R760xs is a rack server appropriate for attended data center environment. However, lower acoustical output is attainable with proper hardware or software configurations.

Configuration	Quietest	12 x 3.5-inch configuration	Quietest - GPU	Volume 1 - 2.5- inch	Volume 2 - 3.5- inch with Rear Storage
CPU TDP	125 W	185 W	150 W	185 W	185 W
CPU Quantity	1	2	2	2	2
RDIMM Memory	16 GB DDR5	32 GB DDR5	16 GB DDR5	16 GB DDR5	32 GB DDR5
Memory Quantity	2	16	4	16	16
Backplane Type	8 x 2.5-inch BP	12 x 3.5-inch BP	8 x 2.5-inch BP	8 x 2.5-inch + 8 x 2.5-inch NVMe BP	12 x 3.5-inch + rear 2 x 2.5-inch BP
HDD Type SATA 2.5-inch 480 GB		SATA 3.5-inch 12 TB	SATA 2.5-inch 480 GB	SAS 2.5-inch + NVMe 2.5-inch	SATA 3.5-inch 12 TB + rear 2.5-inch U.2 NVMe
HDD Quantity	1	12	4	8+8	12+2
Flash Drives	X	X	×	X	X
Flash Quantity	X	X	×	X	X
PSU Type	600 W	1400 W	800 W	800 W	1400 W
PSU Quantity	1	2	2	2	2
OCP	1G	10/25 2-port	1G	10/25 2-port	10/25 2-port
PCI 1	X	X	×	X	X
PCI 2	X	2-port 25Gb	×	2-port 25Gb	2-port 25Gb
PCI 3	X	X	×	X	X
PCI 4	X	X	A2	X	X
PCI 5	X	2-port 25 Gb	×	2-port 25 Gb	2-port 25 Gb
PCI 6	X	X	×	X	X
PERC	Front H355i	Front H755	Front H355i	Front H755	Rear H755

Table 28. Configurations tested for acoustical experience

Table 29. Acoustical experience of R760 configurations

Configuration		Quietest	12 x 3.5-inch configuration	Quietest - GPU	Volume 1 - 2.5- inch	Volume 2 - 3.5-inch with Rear Storage
Acoustical Performance: Idle/ Operating @ 25°C Ambient						
L _{wA,m} (B)	Idle ⁽⁴⁾	4.6	7.0	6.4	6.4	7.2



Table 29. Acoustical experience of R760 configurations (continued)

Configuration		Quietest	12 x 3.5-inch configuration	Quietest - GPU	Volume 1 - 2.5- inch	Volume 2 - 3.5-inch with Rear Storage	
	Operating/Customer usage operating ⁽⁵⁾⁽⁶⁾	4.6	7.0	8.0	6.4	7.2	
К _v (В)	Idle (4)	0.4	0.4	0.4	0.4	0.4	
	Operating/Customer usage operating ⁽⁵⁾⁽⁶⁾	0.4	0.4	0.4	0.4	0.4	
L _{pA,m} (dB)	Idle ⁽⁴⁾	32	56	50	49	57	
	Operating/Customer usage operating ⁽⁵⁾⁽⁶⁾	32	56	64	50	58	
Prominent dis	crete tones ⁽³⁾	Prominence ratio < ECMA-74	Prominence ratio < 15 dB				
Acoustical Pe	rformance: Idle @ 28°C /	Ambient	•				
L _{wA,m} ⁽¹⁾ (B)		4.8	7.2	6.6	6.6	7.4	
К _v (В)		0.4	0.4	0.4	0.4	0.4	
L _{pA,m} ⁽²⁾ (dB)		33	56 52 52 58		58		
Acoustical Performance: Max. loading @ 35°C Ambient							
L _{wA,m} ⁽¹⁾ (B)		5.4 8.0 8.6 8.1 8.1			8.1		
К _v (В)		0.4	0.4	0.4	0.4	0.4	
$L_{pA,m}^{(2)}(dB)$		38	66	70	65	65	

⁽¹⁾LwA,m: The declared mean A-weighted sound power level (LwA) is calculated per section 5.2 of ISO 9296 with data collected using the methods described in ISO 7779 (2010). Engineering data presented here may not be fully compliant with ISO 7779 declaration requirements.

⁽²⁾LpA,m: The declared mean A-weighted emission sound pressure level is at the bystander position per section 5.3 of ISO 9296 and measured using methods described in ISO 7779. The system is placed in a 24U rack enclosure, 25cm above a reflective floor. Engineering data presented here may not be fully compliant with ISO 7779 declaration requirements.

⁽³⁾Prominent tones: Criteria of Annex D of ECMA-74 & Prominence Ratio method of ECMA-418 are followed to determine if discrete tones are prominent and to report them, if so.

⁽⁴⁾Idle mode: Idle mode is the steady-state condition in which the server is energized but not operating any intended function.

⁽⁵⁾Operating mode: Operating mode is represented by the maximum of the steady state acoustical output at 50% of CPU TDP or active storage drives for the respective sections of Annex C of ECMA-74.

⁽⁶⁾ Customer Usage Operating mode: The operating mode is represented by the maximum of the steady state acoustical output at 10%~25% of CPU TDP, 0%~10% IOPs load, and >80% GPU load as the components showed in the above configurations.



Rack, rails, and cable management

Topics:

- Rails information
- Cable Management Arm
- Strain Relief Bar
- Rack Installation

Rails information

The rail offerings for the R760xs consist of two general types: sliding and static. The cable management offerings consist of an optional cable management arm (CMA) and an strain relief bar (SRB).

See the Dell Enterprise Systems Rail Sizing and Rack Compatibility Matrix available at rail-rack-matrix for information regarding:

- Specific details about rail types.
- Rail adjustability ranges for various rack mounting flange types
- Rail depth with and without cable management accessories
- Rack types that are supported for various rack mounting flange types

Key factors governing proper rail selection include the following:

- Identifying the type of rack in which they will be installed.
- The spacing between the front and rear mounting flanges of the rack.
- The type and location of any equipment mounted in the back of the rack such as power distribution units (PDUs), and the overall depth of the rack
- Overall depth of the rack

A11 Sliding Rails features summary

The sliding rails allow the system to be fully extended out of the rack for service. The sliding rails have a Cable Management Arm (CMA) and a Strain Relief Bar (SRB) option.

There are one types of sliding rails available:

• Stab-in/Drop-in sliding rails

A11 Stab-in/Drop-in sliding rails for 4-post racks

- Supports Drop-in or Stab-in installation of the chassis to the rails
- Supports tool-less installation in 19-inch EIA-310-E compliant square, unthreaded round hole racks including all generations of Dell racks.

Also supports tool-less installation in threaded round hole 4-post racks

- Support for tool-less installation in Dell Titan or Titan-D racks
- Supports full extension of the system out of the rack to allow serviceability of key internal components
- (i) **NOTE:** For situations where CMA support is not required, the outer CMA mounting brackets can be uninstalled from the sliding rails. This reduces the overall length of the rails and eliminates the potential interferences with rear mounted PDUs or the rear rack door.

Supports optional Cable Management Arm (CMA)

• Supports optional Strain Relief Bar (SRB)



A8 Static Rails features summary

The static rails, which are shown in the figure below, support a wider variety of racks than the sliding rails, but do not support serviceability in the rack. The static rails are not compatible with the CMA and SRB.

- Supports Stab-in installation of the chassis to the rails
- Supports tool-less installation in 19-inch EIA-310-E compliant square or unthreaded round hole 4-post racks including all generations of Dell racks
- Supports tooled installation in 19-inch EIA-310-E compliant threaded hole 4-post and 2-post racks
- Supports tooled installation in Dell Titan or Titan-D rack

() NOTE:

- Screws are not included with the static rail kit since racks are offered with various thread types.
- Screw head diameter should be 10mm or less.



Figure 30. Static rails

2-Post racks installation

If installing to 2-Post (Telco) racks, the ReadyRails II static rails (A8) must be used. Sliding rails support mounting in 4-post racks only.



Figure 31. Static rails in 2-post center mount configuration

44 Rack, rails, and cable management



Installation in the Dell Titan or Titan-D racks

For tool-less installation in Titan or Titan-D racks, the Stab-in/Drop-in sliding rails (A11) must be used. This rail collapses down sufficiently to fit in the rack with mounting flanges that are spaced about 24 inches apart from front to back. The Stab-in/Drop-in sliding rail allows bezels of the servers and storage systems to be aligned when installed in these racks. For tooled installation, Stab-in Static rails (A8) must be used for bezel alignment with storage systems.

Cable Management Arm

The optional Cable Management Arm (CMA) for the system organizes and secures the cords and cables exiting the back of the server and unfolds to allow the server to extend out of the rack without having to detach the cables.

Some key features of the CMA include:

- Large U-shaped baskets to support dense cable loads
- Open vent pattern for optimal airflow
- Support for mounting on either side by swinging the spring-loaded brackets from one side to the other
- Utilizes hook-and-loop straps rather than plastic tie wraps to eliminate the risk of cable damage during cycling
- Includes a low-profile fixed tray to both support and retain the CMA in its fully closed position
- Both the CMA and the tray mount without the use of tools through simple and intuitive snap-in designs

The CMA can be mounted to either side of the sliding rails without the use of tools or the need for conversion. For systems with one power supply unit (PSU), it is recommended to mount on the side opposite to that of the power supply to allow easier access to it and the rear drives (if applicable) for service or replacement.



Figure 32. Cable Management Arm

Strain Relief Bar

The optional strain relief bar (SRB) for the R760xs organizes and supports cable connections at the rear end of the server to avoid damage from bending.





Figure 33. Cabled strain relief bar

Sliding rails with optional SRB:

- Support tool-less attachment to rails
- Support two depth positions to accommodate various cable loads and rack depths
- Support cable loads and controls stress on server connectors
- Support cables can be segregated into discrete, purpose-specific bundles

Rack Installation

Drop-in design means that the system is installed vertically into the rails by inserting the standoffs on the sides of the system into the J-slots in the inner rail members with the rails in the fully extended position. The recommended method of installation is to first insert the rear standoffs on the system into the rear J-slots on the rails to free up a hand and then rotate the system down into the remaining J-slots while using the free hand to hold the rail against the side of the system.

Stab-in design means that the inner (chassis) rail members must first be attached to the sides of the system and then inserted into the outer (cabinet) members installed in the rack.

Installing system into the rack (option A: Drop-In)

1. Pull the inner rails out of the rack until they lock into place.





Figure 34. Pull out inner rail

- 2. Locate the rear rail standoff on each side of the system and lower them into the rear J-slots on the slide assemblies.
- 3. Rotate the system downward until all the rail standoffs are seated in the J-slots.



Figure 35. Rail standoffs seated in J-slots

- 4. Push the system inward until the lock levers click into place.
- 5. Pull the blue side release lock tabs forward or backward on both rails and slide the system into the rack until the system is in the rack.





Figure 36. Slide system into the rack

Installing the system into the rack (option B: Stab-In)

- 1. Pull the intermediate rails out of the rack until they lock into place.
- 2. Release the inner rail lock by pulling forward on the white tabs and sliding the inner rail out of the intermediate rails.



Figure 37. Pull out the intermediate rail

Table 30. Rail component label

Number	Component
1	Intermediate rail
2	Inner rail

3. Attach the inner rails to the sides of the system by aligning the J-slots on the rail with the standoffs on the system and sliding forward on the system until they lock into place.





Figure 38. Attach the inner rails to the system

4. With the intermediate rails extended, install the system into the extended rails.



Figure 39. Install system into the extended rails

5. Pull blue slide release lock tabs forward or backward on both rails, and slide the system into the rack.





Figure 40. Slide system into the rack

50 Rack, rails, and cable management


Supported operating systems

The PowerEdge R760xs system supports the following operating systems:

- Canonical Ubuntu Server LTS
- Microsoft Windows Server with Hyper-V
- Red Hat Enterprise Linux
- SUSE Linux Enterprise Server
- VMware vSAN/ESXi

For more information, go to www.dell.com/ossupport.

Supported operating systems 51



Dell OpenManage Systems Management

Dell delivers management solutions that help IT administrators effectively deploy, update, monitor, and manage IT assets. OpenManage solutions and tools enable you to quickly respond to problems by helping them to manage Dell servers efficiently; in physical, virtual, local, and remote environments; all without the need to install an agent in the operating system.

The OpenManage portfolio includes:

- Innovative embedded management tools integrated Dell Remote Access Controller (iDRAC)
- Consoles OpenManage Enterprise
- Extensible with plug-ins OpenManage Power Manager
- Update tools Repository Manager

Dell has developed comprehensive systems management solutions that are based on open standards and has integrated with management consoles from partners such as Microsoft and VMware, allowing advanced management of Dell servers. Dell management capabilities extend to offerings from the industry's top systems management vendors and frameworks such as Ansible, Splunk, and ServiceNow. OpenManage tools automate the full span of server life cycle management activities along with powerful RESTful APIs to script or integrate with your choice of frameworks.

For more information about the entire OpenManage portfolio, see:

• The latest Dell Systems Management Overview Guide.

Topics:

- Integrated Dell Remote Access Controller (iDRAC)
- Systems Management software support matrix

Integrated Dell Remote Access Controller (iDRAC)

iDRAC9 delivers advanced, agent-free, local and remote server administration. Embedded in every PowerEdge server, iDRAC9 provides a secure means to automate a multitude of common management tasks. Because iDRAC is embedded within every PowerEdge server, there is no additional software to install; just plug in power and network cables, and iDRAC is ready to go. Even before installing an operating system (operating system) or hypervisor, IT administrators have a complete set of server management features at their fingertips.

With iDRAC9 in-place across the Dell PowerEdge portfolio, the same IT administration techniques and tools can be applied throughout. This consistent management platform allows easy scaling of PowerEdge servers as an organization's infrastructure grows. Customers can use the iDRAC RESTful API for the latest in scalable administration methods of PowerEdge servers. With this API, iDRAC enables support for the Redfish standard and enhances it with Dell extensions to optimize at-scale management of PowerEdge servers. By having iDRAC at the core, the entire OpenManage portfolio of Systems Management tools allows every customer to tailor an effective, affordable solution for any size environment.

Zero Touch Provisioning (ZTP) is embedded in iDRAC. ZTP - Zero Touch Provisioning is Intelligent Automation Dell's agent-free management puts IT administrators in control. Once a PowerEdge server is connected to power and networking, that system can be monitored and fully managed, whether you're standing in front of the server or remotely over a network. In fact, with no need for software agents, an IT administrator can: • Monitor • Manage • Update • Troubleshoot and remediate Dell servers With features like zero-touch deployment and provisioning, iDRAC Group Manager, and System Lockdown, iDRAC9 is purpose-built to make server administration quick and easy. For those customers whose existing management platform utilizes in-band management, Dell does provide iDRAC Service Module, a lightweight service that can interact with both iDRAC9 and the host operating system to support legacy management platforms.

When ordered with DHCP enabled from the factory, PowerEdge servers can be automatically configured when they are initially powered up and connected to your network. This process uses profile-based configurations that ensure each server is configured per your specifications. This feature requires an iDRAC Enterprise license.

iDRAC9 offers following license tiers:



Table 31. iDRAC9 license tiers

License	Description	
iDRAC9 Basic	 Available only on 100-500 series rack/tower Basic instrumentation with iDRAC web UI For cost conscious customers that see limited value in management 	
iDRAC9 Express	 Default on 600+ series rack/tower, modular, and XR series Includes all features of Basic Expanded remote management and server life-cycle features 	
iDRAC9 Enterprise	 Available as an upsell on all servers Includes all features of Basic and Express. Includes key features such as virtual console, AD/LDAP support, and more Remote presence features with advanced, Enterprise-class, management capabilities 	
iDRAC9 Datacenter	 Available as an upsell on all servers Includes all features of Basic, Express, and Enterprise. Includes key features such as telemetry streaming, Thermal Manage, automated certificate management, and more Extended remote insight into server details, focused on high end server options, granular power, and thermal management 	

For a full list of iDRAC features by license tier, see Integrated Dell Remote Access Controller 9 User's Guide at Dell.com. For more details on iDRAC9 including white papers and videos, see:

• Support for Integrated Dell Remote Access Controller 9 (iDRAC9) on the Knowledge Base page at Dell.com

Systems Management software support matrix

Table 32. Systems Management software support matrix

Categories	Features	PE mainstream
Embedded Management and In-band	iDRAC9 (Express, Enterprise, and Datacenter licenses)	Supported
Services	OpenManage Mobile	Supported
	OM Server Administrator (OMSA)	Supported
	iDRAC Service Module (iSM)	Supported
	Driver Pack	Supported
Change Management	Update Tools (Repository Manager, DSU, Catalogs)	Supported
	Server Update Utility	Supported
	Lifecycle Controller Driver Pack	Supported
	Bootable ISO	Supported
Console and Plug-ins	OpenManage Enterprise	Supported
	Power Manager Plug-in	Supported
	Update Manager Plug-in	Supported
	SupportAssist Plug-in	Supported
	CloudIQ	Supported
Integrations and connections	OM Integration with VMware Vcenter/vROps	Supported
	OM Integration with Microsoft System Center (OMIMSC)	Supported
	Integrations with Microsoft System Center and Windows Admin Center (WAC)	Supported



Table 32. Systems Management software support matrix (continued)

Categories	Features	PE mainstream
	ServiceNow	Supported
	Ansible	Supported
	Third-party Connectors (Nagios, Tivoli, Microfocus)	Supported
Security	Secure Enterprise Key Management	Supported
	Secure Component Verification	Supported
Standard operating system	Red Hat Enterprise Linux, SUSE, Windows Server 2019 or 2022, Ubuntu, CentOS	Supported (Tier-1)

54 Dell OpenManage Systems Management



Appendix A. Standards compliance

The system conforms to the following industry standards.

Table 33. Industry standard documents

Standard	URL for information and specifications	
ACPI Advance Configuration and Power Interface Specification, v6.4	Uefi specifications and tools	
Ethernet IEEE Std 802.3-2022	ieee standards	
MSFT WHQL Microsoft Windows Hardware Quality Labs	microsoft.com/whdc/system/platform/pcdesign/desguide/ serverdg.mspx	
IPMI Intelligent Platform Management Interface, v2.0	intel.com/design/servers/ipmi	
DDR5 Memory DDR5 SDRAM Specification	jedec.org/standards-documents/docs/jesd79-4.pdf	
PCI Express PCI Express Base Specification, v5.0	pcisig.com/specifications/pciexpress	
PMBus Power System Management Protocol Specification, v1.2	pmbus specification and revisions	
SAS Serial Attached SCSI, 3 (SAS-3) (T10/INCITS 519)	SCSI storage interfaces information	
SATA Serial ATA Rev. 3.3	sata-io.org page	
SMBIOS System Management BIOS Reference Specification, v3.3.0	BIOS reference specification page	
TPM Trusted Platform Module Specification, v1.2 and v2.0	trustedcomputinggroup org page	
UEFI Unified Extensible Firmware Interface Specification, v2.7	UEFIF specifications	
PI Platform Initialization Specification, v1.7		
USB Universal Serial Bus v2.0 and SuperSpeed v3.0 (USB 3.1 Gen1)	USB Implementers Forum, Inc. USB document library	
NVMe Express Base Specification. Revision 2.0c	NVME specifications	
 NVMe Command Set Specifications NVM Express NVM Command Set Specification. Revision 1.1c NVM Express Zoned Namespaces Command Set. Revision 1.0c NVM Express® Key Value Command Set. Revision 1.0c 		
 NVMe Transport Specifications NVM Express over PCle Transport. Revision 1.0c NVM Express RDMA Transport Revision. 1.0b NVM Express TCP Transport. Revision 1.0c 		
NVMe NVM Express Management Interface. Revision 1.2c		
NVMe NVMe Boot Specification. Revision 1.0		



Appendix B: Additional specifications

Topics:

- Chassis dimensions
- System weight
- Video specifications
- USB ports specifications
- Environmental specifications

Chassis dimensions



Figure 41. Chassis dimensions

Table 34. PowerEdge R760xs chassis dimensions

Xa	Xb	Y	Za	Zb	Zc
482.0 mm (18.97 inches)	434.0 mm (17.08 inches)	86.8 mm (3.41 inches)	22.0 mm (0.86 inches) Without bezel	677.44 mm (26.67 inches) Ear to L bracket housing	685.78 mm (26.99 inches) Ear to PSU handle without velcro strap

56 Appendix B: Additional specifications

Inserido ao protocolo **22.951.206-4** por: **Pedro Henrique Golin Linhares** em: 31/03/2025 12:49. A autenticidade deste documento pode ser validada no endereço: https://www.eprotocolo.pr.gov.br/spiweb/validarDocumento com o código: fc1c44dfe0686d5547ffd691199be699.



Table 34. PowerEdge R760xs chassis dimensions

Xa	ХЬ	Y	Za	Zb	Zc
			35.84 mm (1.41 inches) With bezel	650.24 mm (25.6 inches) Ear to PSU surface	

(i) NOTE: Zb is the nominal rear wall external surface where the system board I/O connectors reside.

System weight

Table 35. PowerEdge R760xs systemweight

System configuration	Maximum weight (with all drives/SSDs/bezel)
16+8 x 2.5-inch	25.92 kg (57.14 lb)
16 x 2.5-inch	24.58 kg (54.18 lb)
12 x 3.5-inch	28.82 kg (63.53 lb)
8 x 3.5-inch	25.84 kg (54.96 lb)
8 x 2.5-inch	21.56 kg (47.53 lb)
No backplane configuration	19.40 kg (42.76 lb)

Video specifications

The PowerEdge R760xs system supports integrated Matrox G200 graphics controller with 16 MB of video frame buffer.

Table 36. Supported video resolution options for the system

Resolution	Refresh rate (Hz)	Color depth (bits)
1024 x 768	60	8, 16, 32
1280 x 800	60	8, 16, 32
1280 x 1024	60	8, 16, 32
1360 x 768	60	8, 16, 32
1440 x 900	60	8, 16, 32
1600 x 900	60	8, 16, 32
1600 x 1200	60	8, 16, 32
1680 x 1050	60	8, 16, 32
1920 x 1080	60	8, 16, 32
1920 x 1200	60	8, 16, 32



USB ports specifications

Table 37. PowerEdge R760xs USB specifications

Front		Rear		Internal (Optional)	
USB port type	No. of ports	USB port type	No. of ports	USB port type	No. of ports
USB 2.0- compliant port	One	USB 2.0- compliant port	One	Internal USB 3.0- compliant port	One
iDRAC Direct port (Micro-AB USB 2.0- compliant port)	One	USB 3.0- compliant port	One		

(i) NOTE: The micro USB 2.0 compliant port can only be used as an iDRAC Direct or a management port.

NOTE: The USB 2.0 specifications provide a 5 V supply on a single wire to power connected USB devices. A unit load is defined as 100 mA in USB 2.0, and 150 mA in USB 3.0. A device may draw a maximum of 5 unit loads (500 mA) from a port in USB 2.0; 6 (900 mA) in USB 3.0.

NOTE: The USB 2.0 interface can provide power to low-power peripherals but must adhere to USB specification. An external power source is required for higher-power peripherals to function, such as external CD/DVD Drives.

Environmental specifications

NOTE: For additional information about environmental certifications, see the Product Environmental Datasheet that are located with the Documentation on https://www.dell.com/support.

Table 38. Continuous Operation Specifications for ASHRAE A2

Temperature, humidity and, operational altitude	Allowable continuous operations
Temperature range for altitudes <= 900 m (<= 2953 ft)	10-35°C (50-95°F) with no direct sunlight on the equipment
Humidity percent range (non-condensing at all times)	8% RH with -12°C (10.4°F) minimum dew point to 80% RH with 21°C (69.8°F) maximum dew point
Operational altitude derating	Maximum temperature is reduced by 1°C/300 m (1.8°F/984 Ft) above 900 m (2953 Ft)

Table 39. Continuous Operation Specifications for ASHRAE A3

Temperature, humidity and, operational altitude	Allowable continuous operations
Temperature range for altitudes <= 900 m (<= 2953 ft)	5–40°C (41–104°F) with no direct sunlight on the equipment
Humidity percent range (non-condensing at all times)	8% RH with -12°C (10.4°F) minimum dew point to 85% RH with 24°C (75.2°F) maximum dew point
Operational altitude derating	Maximum temperature is reduced by 1°C/175 m (1.8°F/574 Ft) above 900 m (2953 Ft)

Table 40. Continuous Operation Specifications for ASHRAE A4

Temperature, humidity and, operational altitude	Allowable continuous operations
Temperature range for altitudes <= 900 m (<= 2953 ft)	5-45°C (41-113°F) with no direct sunlight on the equipment

58 Appendix B: Additional specifications



Table 40. Continuous Operation Specifications for ASHRAE A4 (continued)

Temperature, humidity and, operational altitude	Allowable continuous operations		
Humidity percent range (non-condensing at all times)	8% RH with -12°C (10.4°F) minimum dew point to 90% RH with 24°C (75.2°F) maximum dew point		
Operational altitude derating	Maximum temperature is reduced by 1°C/125 m (1.8°F/410 Ft) above 900 m (2953 Ft)		

Table 41. Continuous Operation Specifications for Rugged Environment

Temperature, humidity and, operational altitude	Allowable continuous operations
Temperature range for altitudes <= 900 m (<= 2953 ft)	5-45°C (41-113°F) with no direct sunlight on the equipment
Humidity percent range (non-condensing at all times)	8% RH with -12°C (10.4°F) minimum dew point to 90% RH with 24°C (75.2°F) maximum dew point
Operational altitude derating	Maximum temperature is reduced by 1°C/125 m (1.8°F/410 Ft) above 900 m (2953 Ft)

Table 42. Common Environmental Specifications for ASHRAE A2, A3, A4, and Rugged

Temperature, humidity and, operational altitude	Allowable continuous operations
Maximum temperature gradient (applies to both operation and non-operation)	20°C in an hour* (36°F in an hour) and 5°C in 15 minutes (9°F in 15 minutes), 5°C in an hour* (9°F in an hour) for tape (i) NOTE: * - Per ASHRAE thermal guidelines for tape hardware, these are not instantaneous rates of temperature change.
Non-operational temperature limits	-40°C to 65°C (-104°F to 149°F)
Non-operational humidity limits	5% to 95% RH with 27°C (80.6°F) maximum dew point
Maximum non-operational altitude	12,000 meters (39,370 feet)
Maximum operational altitude	3,048 meters (10,000 feet)

Table 43. Maximum vibration specifications

Maximum vibration	Specifications		
Operating	0.21 G_{rms} at 5 Hz to 500 Hz for 10 minutes (all operation orientations)		
Storage	1.88 G_{rms} at 10 Hz to 500 Hz for 15 minutes (all six sides tested)		

Table 44. Maximum shock pulse specifications

Maximum shock pulse	Specifications			
Operating	Six consecutively performed shock pulses in the positive and negative x, y, and z axis of 6 G for up to 11 ms.			
Storage	Six consecutively performed shock pulses in the positive and negative x, y, and z axis (one pulse on each side of the system) of 71 G for up to 2 ms.			

The following table defines the limitations that help avoid any equipment damage or failure from particulates and gaseous contamination. If the levels of particulate or gaseous pollution exceed the specified limitations and result in equipment damage or failure, you may need to rectify the environmental conditions. Re-mediation of environmental conditions is the responsibility of the customer.



Table 45. Particulate and gaseous contamination specifications

Particulate contamination	Specifications			
Air filtration	 Data center air filtration as defined by ISO Class 8 per ISO 14644-1 with a 95% upper confidence limit. (i) NOTE: The ISO Class 8 condition applies to data center environments only. This air filtration requirement does not apply to IT equipment designed to be used outside a data center, in environments such as an office or factory floor. (i) NOTE: Air entering the data center must have MERV11 or MERV13 filtration. 			
Conductive dust	Air must be free of conductive dust zinc whiskers, or other conductive particles. (i) NOTE: This condition applies to data center and non-data center environments.			
Corrosive dust	Air must be free of corrosive dust. Residual dust present in the air must have a deliquescent point less than 60% relative humidity. (i) NOTE: This condition applies to data center and non-data center environments.			

Table 46. Gaseous contamination specifications

Gaseous contamination	Specification		
Copper coupon corrosion rate	<300A/month per class G1 as defines by ANSI/ISA71.04-2013		
Silver coupon corrosion rate	<200A/month as defined by ANSI/ISA71.04-2013		

Thermal restriction matrix

Table 47. Processor and heat sink matrix

Heat sink	Processor TDP	
STD HSK	< 185 W CPU SKUs	
HPR HSK	185 W-250 W CPU SKUs (12 x 3.5-inch drive configuration not supported)	
	125W-250W CPU SKUs (12 x 3.5-inch drive configuration supported)	

Table 48. Label reference

Label	Description		
STD	Standard		
HPR (Silver)	High Performance Silver (HPR) fan		
HPR (Gold)	High Performance Gold (VHP) fan		
HSK	Heat sink		

() NOTE: The ambient temperature of the configuration is determined by the critical component in that configuration. For example, if the processor's supported ambient temperature is 35°C, the DIMM is 35°C, and the GPU is 30°C, the combined configuration can only support 30°C.

60 Appendix B: Additional specifications



Table 49. Supported ambient temperature for processors for R760xs

R760xs										
configur	ation		No backpla ne	8 x 3.5- inch SAS configu ration	12 x 3.5- inch SAS configu ration	12 x 3.5- inch configu ration with rear drive module	8 x 2.5- inch SAS configu ration	8 x 2.5- inch NVMe configu ration	16 x 2.5- inch SAS configu ration	16 x 2.5- inch + 8 x 2.5- inch NVMe configu ration
EMR	4514Y	150 W	45°C	40°C	35°C	35°C	40°C	40°C	40°C	40°C
MCC CPU	5512U	185 W	45°C	40°C	35°C	35°C	40°C	40°C	40°C	40°C
	6526Y/6534	195 W	35°C	35°C	35°C	35°C	35°C	35°C	35°C	35°C
	6548Y+/ 6542Y/6548N	250 W	35°C	35°C	30°C	30°C	35°C	35°C	35°C	35°C
SPR	4509Y	125 W	45°C	40°C	35°C	35°C	40°C	40°C	40°C	40°C
CPU	4510	150 W	45°C	40°C	35°C	35°C	40°C	40°C	40°C	40°C
SPR	3408U	125 W	45°C	40°C	35°C	35°C	40°C	40°C	40°C	40°C
MCC CPU	5416S/ 4410T/ 4410Y/5415+	150 W	45°C	40°C	35°C	35°C	40°C	40°C	40°C	40°C
	4416	165 W	45°C	40°C	35°C	35°C	40°C	40°C	40°C	40°C
	5418Y/ 5412U/6426Y	185 W	45°C	40°C	35°C	35°C	40°C	40°C	40°C	40°C
	5420+/ 6438Y+	205 W	35°C	35°C	35°C	35°C	35°C	35°C	35°C	35°C
	6448Y/6442Y	225 W	35°C	35°C	35°C	35°C	35°C	35°C	35°C	35°C
	6414U	250 W	35°C	35°C	30°C	30°C	35°C	35°C	35°C	35°C
Memory	96 GB RDIMM 5200	8.1 W, 1DPC	35°C	35°C	30°C	30°C	35°C	35°C	35°C	35°C
	64 GB RDIMM 5200	7.7 W, 1DPC	45°C	40°C	35°C	35°C	40°C	40°C	40°C	40°C
	32 GB RDIMM 5200	5.1 W, 1DPC	45°C	40°C	35°C	35°C	40°C	40°C	40°C	40°C
	64 GB RDIMM 4800	12 W, 1DPC	45°C	40°C	35°C	35°C	40°C	40°C	40°C	40°C
	32 GB RDIMM 4800	10 W, 1DPC	45°C	40°C	35°C	35°C	40°C	40°C	40°C	40°C
PCle		45°C	40°C	35°C1	35°C ¹	40°C	40°C	40°C	40°C	
A2 GPU ⁶		35°C	30°C	Not supporte d	Not supporte d	35°C	35°C	30°C	30°C	
OCP			45°C	40°C	35°C ²	35°C ²	40°C	40°C	40°C	40°C
BOSS			35°C	35°C	35°C	35°C	35°C	35°C	35°C	35°C

(i) NOTE:

- 1. Max supported thermal tier of PCIe card is Tier 5.
- **2.** Max supported thermal tier of OCP is Tier 5.



- HPR Sliver fan is required from fan zone 2 to fan zone 6 for 8 x 2.5-inch NVMe, 16 x 2.5-inch SAS/SATA + 8 x 2.5-inch NVMe, 12 x 3.5-inch drives or GPU configurations.
- 4. Optional fan zone 1 has to be populated with HPR Gold fan is for BOSS, GPU or rear drive module populations.
- 5. PCIe slot priority of Nvidia A2 GPU is constrained on slot #3, #4, #6.
- 6. HPR heatsink is required for ≥ 185 W CPUs, 12 x 3.5-inch drives or 12 x 3.5-inch drives with rear storage module configurations.
- 7. DIMM blank is required for 12 x 3.5-inch SAS/SATA with rear storage module.
- 8. Fan blank is required on fan zone 1 when no fan population.
- **9.** OCP shroud is required for OCP card population without PCIe riser module installed.
- 10. CPU blank is required for single processor configuration.
- 11. Rear drive module does not support Kioxia CM6 series, Samsung PM1735 series, Hynix PE8010/PE8110 ≥ 7.68 TB,
- Samsung PM1733a > 1.92 TB, Samsung PM1735a > 1.6 TB and Redtail NVMe drive.

(i) NOTE: The fan speed in the 3.5-inch chassis is limited to 90% due to the drive dynamic profile.

Table 50. Fan population rule for R760xs

configuratio n	No backplane	8 x 3.5-inch SAS	12 x 3.5- inch SAS	8 x 2.5-inch SAS	8 x 2.5-inch NVMe	16 x 2.5-inch SAS	24 x 2.5-inch (16 x 2.5-inch
Optional HW							NVMe)
Default	Fan 2 to Fan 6 with STD fan	Fan 2 to Fan 6 with STD fan	Fan 2 to Fan 6 with HPR Silver fan	Fan 2 to Fan 6 with STD fan	Fan 2 to Fan 6 with HPR Silver fan	Fan 2 to Fan 6 with STD fan	Fan 2 to Fan 6 with HPR Silver fan
Rear Module	Not supported	Not supported	Fan 1 with HPR Gold fan	Not supported	Not supported	Not supported	Not supported
			Fan 2 to Fan 6 with HPR Silver fan				
BOSS N1	Fan 1 with HPR Gold fan	Fan 1 with HPR Gold fan	Fan 1 with HPR Gold fan				
	Fan 2 to Fan 6 with STD fan	Fan 2 to Fan 6 with STD fan	Fan 2 to Fan 6 with HPR Silver fan	Fan 2 to Fan 6 with STD fan	Fan 2 to Fan 6 with HPR Silver fan	Fan 2 to Fan 6 with STD fan	Fan 2 to Fan 6 with HPR Silver fan
GPU	Fan 1 with HPR Gold fan	Fan 1 with HPR Gold fan	Not supported	Fan 1 with HPR Gold fan	Fan 1 with HPR Gold fan	Fan 1 with HPR Gold fan	Fan 1 with HPR Gold fan
	Fan 2 to Fan 6 with HPR Silver fan	Fan 2 to Fan 6 with HPR Silver fan		Fan 2 to Fan 6 with HPR Silver fan	Fan 2 to Fan 6 with HPR Silver fan	Fan 2 to Fan 6 with HPR Silver fan	Fan 2 to Fan 6 with HPR Silver fan

Thermal Restrictions for PCIe adapter NIC and other network cards with iDRAC

- Cannot support PCIe cards with the cooling requirement more than 300LFM at 55C in a 12 x 3.5-inch SAS/SATA configuration.
- Solarflare Melrose DP 25 GBE SFP28 (TTKWY) not supported with 12 x 3.5-inch SAS/SATA configuration.
- 100 Gb network adapter or 100 Gb OCP is not supported in the 12 x 3.5-inch SAS/SATA configuration.
- Few 25 GB OCP cards with the cooling requirement more than 250LFM at 55C (3Y64D/4TRD3 / GGGDF/R1KTR / Y4VV5) is not supported in 12x3.5" SAS/SATA configuration.
- The 12 x 3.5-inch SAS/SATA configuration requires the optical transceiver with higher temperature spec (\ge 85°C) to support (M14MK / N8TDR).



- Quad port OCP (3Y64D/Y4VV5) requires the optical transceiver with higher temperature spec (≥ 85°C) to support (M14MK).
- 100 Gb network adapter cannot support the transceivers as 14NV5/9JKK2 / QSFP56 (MFS1S00-VxxxE/HxxxE).
- The H965e is limited to populate in PCI slot 3 in a 12 x 3.5-inch SAS/SATA configuration.
- Mellanox CX7 NDR200 card has few limitations of PCI slot locations.

Table 51. Mellanox CX7 NDR200 slot location limitations

Storage configuration	Slots on 3.5-inch configuration	Slots on 2.5-inch configuration
Gen5 PCIe sloit support for CX7 NDR200	3, 4	3, 4
Gen4 PCle sloit support for CX7 NDR200	6	5, 6

Thermal restrictions for extended ambient support (ASHRAE A3/A4)

- Two PSUs are required in redundant mode. Single PSU failure is not supported.
- 12 x 3.5-inch SAS/SATA configuration is not supported.
- BOSS(M.2) module is not supported.
- CPU TDP > 185 W is not supported.
- PCle card TDP > 25 W is not supported.
- OCP cards with transmission rate higher than 25 GB is not supported.
- OCP transceiver spec \leq 75°C is not supported.
- 8 x 3.5-inch SAS/SATA, 8 x 2.5-inch SAS/SATA, 8 x 2.5-inch NVMe, 16 x 2.5-inch SAS/SATA, 16 x 2.5-inch SAS/SATA + 8x 2.5-inch NVMe configurations are limited to support A3.
- 128 GB+ memory is not supported .
- The rear drive is not supported.



Appendix C Additional resources

Table 52. Additional resources

Resource	Description of contents	Location
Installation and Service Manual	This manual, available in PDF format, provides the following information:	Dell.com/Support/Manuals
	 Chassis features System Setup program System indicator codes System BIOS Remove and replace procedures 	
	DiagnosticsJumpers and connectors	
Getting Started Guide	This guide ships with the system, and is also available in PDF format. This guide provides the following information:	Dell.com/Support/Manuals
	Initial setup steps	
Rack Installation Guide	This document ships with the rack kits, and provides instructions for installing a server in a rack.	Dell.com/Support/Manuals
System Information Label	The system information label documents the system board layout and system jumper settings. Text is minimized due to space limitations and translation considerations. The label size is standardized across platforms.	Inside the system chassis cover
QR code for system resources	This code on the chassis can be scanned by a phone application to access additional information and resources for the server, including videos, reference materials, service tag information, and Dell contact information.	Inside the system chassis cover
Enterprise Infrastructure Planning Tool (EIPT)	The Dell online EIPT enables easier and more meaningful estimates to help you determine the most efficient configuration possible. Use EIPT to calculate the power consumption of your hardware, power infrastructure, and storage.	Dell.com/calc

Topics:

- Documentation
- Customer kits



Documentation

This section provides information about the documentation resources for your system.

Table 53. Documentation resources

Document	Location
Factory Configuration Matrix	Sales Portal
SPM (Slot Priority Matrix)	Sales Portal
NDA Deck	Sales Portal
Installation and Service Manual (ISM)	https://www.dell.com/poweredgemanuals
Field Service Manual (FSM)	https://www.dell.com/poweredgemanuals > Sing in
Technical Guide	Dell.com > Product page > Product Details
Spec Sheet	Dell.com > Product page > Product Details

Customer kits

Dell Upgrades

It is not always possible to plan for new applications, future workloads, and business needs. Unleash the full power of your Dell Technologies Infrastructure. When budget does not permit the purchase of new servers, Dell Upgrades is a cost-effective method to repurpose and unleash the full power of existing server, storage, and networking infrastructure.

- Protect your mission-critical operations by using only genuine Dell OEM-validated Upgrades and the technical expertise of Dell ProSupport
- Flex and scale existing infrastructure by upgrading, adding memory or storage drives to cost-effectively and quickly meet new workloads and demands
- Dell Upgrades are the same peripheral commodities that your customer may improve or maintain their server after the initial point of sale

Upgrades portfolio

Table 54. Upgrade category

Dell Upgrade Category	Sample Picture	Dell Upgrade Category Offerings
Memory Memory upgrades are essential for keeping your customers operating at peak performance as their business needs grow and their workloads increase. We tend to see strong demand for server memory because it is the easiest and most cost-effective way to improve system performance.	Contraction of the second seco	DDR5 5600 MT/s and 4800 MT/s



Table 54. Upgrade category (continued)

Dell Upgrade Category	Sample Picture	Dell Upgrade Category Offerings
Storage Dell offers both solid-state drive and hard disk drive storage options for enterprise systems with SATA, SAS or NVMe interfaces. SSDs excel in speed, high-performance I/O requirements, and high reliability due to the lack of spinning disks. Hard Disk Drives (HDDs) store data on spinning disks and offer value for the amount of data storage for the price. Dell offers both solid-state drive and hard disk drive storage options for enterprise systems with SATA, SAS interfaces. SSDs excel in speed, high- performance I/O requirements, and high reliability due to the lack of spinning disks. Hard Disk Drives (HDDs) store data on spinning disks and offer value for the amount of data storage for the price.		HDD: SATA, SAS interface SSD: SATA, SAS, PCI NVMe interface Tape Drive or Media
Processor Processor upgrades help customers perform and accomplish more tasks overall, saving them valuable time. Our processor upgrades include Intel® Xeon® Scalable processors to meet your customers workload needs with increased cores and improved security.		Processors (Intel) Heat sinks
Networking and Optics Our networking and optics components —network interface cards, transceivers, optical cables, and more—are key in today's data center environment, helping customers to improve bandwidth to better manage increase in workloads, devices, users, and interconnected systems.		Network cards Transceivers (Optics)



Table 54. Upgrade category (continued)

Dell Upgrade Category	Sample Picture	Dell Upgrade Category Offerings
Accessories: Dell sells accessories like	8	Controller cards
bezels, controller cards, GPU, PERC and	1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1	Power supplies
other components to complete the Dell Upgrades portfolio and redundancies.	American and	Cables
	· .	Rail kits
		Bezels
	-	Power cords
	T	GPU
	Delland The	PERC
	[80]	BOSS
		Power cords
	States	Cable Management Arm (CAM)
		Fans
		Serial board
	NOT OCC	Internal USB
	T contraction of the second se	
	2800W	

Upgrades reference links

- Main Upgrades Page
- Customer Kit Selector
- Dell Parts Finder Tool (Customer Facing Tool)



Appendix D: Service and support

Topics:

- Why attach service contracts
- ProSupport Infrastructure Suite
- Specialty Support Services
- Dell deployment services
- Supplemental Deployment Services
- Unique Deployment Scenarios
- DAY 2 Automation Services with Ansible
- Dell Technologies Consulting Services

Why attach service contracts

Dell PowerEdge servers include a standard hardware warranty that highlights our commitment to product quality by guaranteeing repair or replacement of defective components. While industry-leading, our warranties are limited to 1 or 3 years, depending on model, and do not cover software assistance. Call records show that failure rates for servers are roughly 1% and more commonly, customers seek Dell technical support for software-related issues like configuration guidance, troubleshooting, upgrade assistance, or performance tuning. Encourage customers to purchase ProSupport service contracts to supplement warranty coverage and ensure optimal support for both hardware and software. ProSupport provides a complete hardware guarantee beyond the original warranty period (up to 12 years: including seven years standard support and an additional five years of Post-Standard Support). Details of the ProSupport Suite and benefits are listed below.

ProSupport Infrastructure Suite

ProSupport Infrastructure Suite is a set of support services that enable customers to build the solution that is right for their organization. It is an industry-leading, enterprise-class support that aligns with the criticality of your systems, the complexity of your environment, and the allocation of your IT resources.



ProSupport Infrastructure Suite | Enhanced value across all offers!

	Basic Herdware Support	ProSupport For Intrastructure	ProSupport Plus for Infrastructure	Changes with August 2025 minute
Technical support availability and response objective	Why intermediated	247, immediate	24/7, mmodiate	No shange
Covered producte	Hintkeara	Hantwere & Software	Hartheare & Software	No shanga
Onsite response service level	1000	NSO or 4 hour	4-hour	Projugator Plus NBD is relived.
ProSupport AlOps platforms			•	Ny Service 010 and TechDirect (all offers) Close00 (Perdisport & ProSupport Plus)
Del Security Advisories	(III)		•	Available on additional products
Proactive issue detection with automated case creation		•	•	New to Searc
Predictive hardware anomaly detection			•	New In ProSupport
Access to software updates		<u> </u>	•	No change
GoudiO health and cybersecurity monitoring & analytics		100 C	•	Estances Nervine
Incident Manager for Severity 1 cases		1 Qu	•	No change
Mission Onlinal support			•	Elifiances features
Priority access to remote senior support engineers1				No change
Service Account Manager			•	No uhonge
Proactive system maintenance				No charge
Limited 3 rd party software support ^{er}		1		No churge

*Software looses can be pumbased through Dell or BYOL - see Service Descriptions for details

DELL'actuations

Figure 42. ProSupport Enterprise Suite

ProSupport Plus for Infrastructure

ProSupport Plus for Infrastructure is the ultimate solution for customers seeking preventative maintenance and optimal performance on their business-critical assets. The service caters to customers who require proactive, predictive, and personalized support for systems that manage critical business applications and workloads. When customers purchase PowerEdge server, we recommend ProSupport Plus, our proactive and preventative support service for business-critical systems. ProSupport Plus provides all the benefits of ProSupport, including the following "Top five reasons to buy ProSupport Plus (PSP)"

- 1. **Priority access to specialized support experts:** Immediate, advanced troubleshooting from an engineer that understands Dell infrastructure solutions.
- 2. Mission Critical Support: When critical (Severity 1) support issues happen, the customer is assured that we do all that we can to get them back up and running as quickly as possible.
- 3. Service Account Manager: A customer's #1 support advocate, ensuring they get the best possible proactive and predictive support experience.
- 4. Systems maintenance: On a semiannual basis, we will keep a customer's ProSupport Plus system(s) up to date by installing the latest firmware, BIOS, and driver updates to improve performance and availability.
- 5. Third-party software support: Dell is a customer's single point of accountability for any eligible third-partysoftware that is installed on their ProSupport Plus system, whether they purchased the software from us or not.

ProSupport for Infrastructure

Comprehensive 24x7 support for hardware and software – best for production, but not critical, workloads and applications. The ProSupport service offers highly trained experts around the clock and around the globe to address IT needs. We help minimize disruptions and maximize availability of PowerEdge server workloads with:

- 24x7 support through phone, chat and online
- A central point of accountability for all hardware and software issues
- Hypervisor, operating system and application support
- Dell security advisories
- Onsite response service levels 4 hour or Next Business Day options
- Proactive issue detection with automated case creation



- Predictive hardware anomaly detection
- Incident Manager assigned for Severity 1 cases
- Collaborative third-party support
- Access to AIOps Platforms (MyService360, TechDirect, and CloudIQ)
- Consistent experience regardless of where customers are located or what language that they speak.

Basic Hardware Support

Provides reactive hardware support during normal business hours, excluding local national holidays. No software support orsoftware-related guidance. For improved levels of support, choose ProSupport or ProSupport Plus.

Specialty Support Services

Optional specialty support services complement the ProSupport Infrastructure Suite to provide additional proficiencies that are critical for modern data center operations.

Hardware coverage add-ons to ProSupport

• Keep Your Hard Drive (KYHD), Keep Your Component (KYC), or Keep Your GPU:

Normally if a device fails under warranty, Dell replaces it using a one-for-one exchange process.KYHD/KYCC/KYGPU gives you the option to retain your device. It provides full control of sensitive data and minimizes security risk by letting you retain possession of failed drives, components, or GPU when receiving replacement parts without incurring additional cost.

• Onsite Diagnosis Service:

Ideal for sites with non-technical staff. Dell field technician performs initial troubleshooting diagnosis onsite and transfers to Dell remote engineers to resolve the issue.

ProSupport Add-on for HPC:

Sold as an add-on to a ProSupport service contract, the ProSupport Add-on for HPC provides solution-aware support to cover the additional requirements that are required to maintain an HPC environment such as:

- \circ $\,$ Access to senior HPC experts
- Advanced HPC cluster assistance: performance, interoperability, and configuration
- Enhanced HPC solution level end-to-end support
- Remote pre-support engagement with HPC Specialists during ProDeploy implementation
- ProSupport Add-on for Telco (Respond & Restore):

An add-on service designed for the top 31 TELCO customers globally, Respond & Restore provides direct access to Dell solution experts who specialize in TELCO carrier-grade support. This add-on also provides a hardware uptime guarantee, meaning if a system fails, Dell has it installed and operational within 4 hours for Severity 1 issues. Dell incurs penalties and fees if SLAs are not met.

Personalized Support and Supplemental Site-wide Expertise

Technical Account Manager:

Designated technology lead who monitors and manages the performance and configuration of specific technology sets.

• Designated Remote Support:

Personalized support expert who manages all troubleshooting and resolution of IT assets.

• Multivendor Support Service:

Support your third-party devices as one service plan for servers, storage, and networking (includes coverage for: Broadcom, Cisco, Fujitsu, HPE, Hitachi, Huawei, IBM, Lenovo, NetApp, Oracle, Quanta, SuperMicro and others).



Services for large enterprises

• ProSupport One for Data Center:

ProSupport One for Data Center offers flexible site-wide support for large and distributed data centers with more than 1,000 assets (combined total of server, storage, networking, so forth). This offering is built on standard ProSupport features that leverage our global scale and are tailored to specific customer needs. While not for everyone, this service option offers a truly unique solution for our largest customers with the most complex environments.

- Team of assigned Services Account Managers with remote or onsite options
- Assigned technical and field engineers who are trained on the customer's environment and configurations.
- On-demand reporting and recommendations that are enabled by ProSupport AlOps tools (MyService360, TechDirect, and CloudIQ)
- Flexible onsite support and parts options that fit their operational model
- A tailored support plan and training for their operations staff
- ProSupport One for CSPs (Cloud Serviced Providers)

ProSupport One for CSPs is a unique offer that is designed for a limited set of Dell accounts purchasing Gen Al computing solutions greater than 1,000 servers and \$250M in sales. PS1 for CSPs improves the entire services experience combining support, deployment (rack integration), residency services, a designated support engineer and the LOIS parts locker as one holistic bundle. Special pricing has been determined to compete effectively against competitors and provide the best customer experience. PS1 for CSPs can only be sold with XE Servers and all networking platforms (Dell and NVIDIA). All other products would be eligible for the standard PS1DC not this special bundle offer. More details on PS1 for CSPs here.

• Logistics Online Inventory Solution (LOIS)

Ideal for large organizations that have their own staff to support their data center. Dell offers a service that is called Logistics Online Inventory Solution which is an onsite parts locker that provides self-maintainers with a local inventory of common replacement components. Having access to these parts lockers allows the self-maintainer to replace a failed component immediately without delay. Each replacement part would automatically initiate a replenishment of the parts inventory that is shipped the next day or delivered onsite by Dell during a regular scheduled visit (called Scheduled Onsite Service). As part of the LOIS system, customers can integrate their systems directly to Dell TechDirect using APIs to help streamline the support management process.

End-of-Life Services

• Post Standard Support (PSS)

Extend service life beyond the initial seven years of ProSupport, adding up to five more additional years of hardware coverage.

Data Sanitization & Data Destruction

Renders data unrecoverable on repurposed or retired products, ensuring security of sensitive data and enabling compliance and provides NIST-compliant certification.

• Asset Recovery Services

Recycle, resale, and disposal of hardware. Helps you securely and responsibly retire IT assets that are no longer needed while protecting both your business and the planet.

Dell deployment services

Dell ProDeploy Infrastructure Suite

ProDeploy Infrastructure Suite provides a variety of deployment offerings that satisfy a customer's unique needs. It is made up of 5 offers: ProDeploy Configuration Services, ProDeploy Rack Integration Services, Basic Deployment, ProDeploy, and ProDeploy Plus.

ProDeploy Infrastructure Suite for servers

Versatile choices for accelerated deployments



Fls.<u>1588</u> Mov. 33

Figure 43. ProDeploy Infrastructure Suite for servers

The new Factory Services consist of two tiers of deployment that happen prior to shipping to the customer's site.

Factory Based Services:

- ProDeploy Factory Configuration Ideal for customers buying servers in volume and seeking pre-configuration prior to shipping such as: custom image, system settings, and asset tagging so it arrives ready to use out of the box. Furthermore, servers can be packaged and bundled to meet specific shipping and distribution requirements for each customer location to facilitate the rollout process. Upsell one of the field based services (below) if a customer needs assistance with the final server installation.
- ProDeploy Rack Integration Ideal for customers seeking to build out fully integrated racks prior to shipping. These rack builds include hardware install, cabling, and full system configuration. You can also add-on a factory stress test and optional on-site final rack configuration to complete the rack installation.
 - STANDARD SKUs for Rack Integration is available in US only and requires:
 - 20 or more devices (R and C series servers and all Dell or non-Dell switches). Use Informational SKUs for Dell switches or 3rd party products
 - Shipping to contiguous US
 - USE CUSTOM QUOTE for Rack Integration for:
 - All countries except USA
 - Racks containing less than 20 servers
 - Any rack that includes VxRail or Storage
 - Shipping outside contiguous US
 - Shipping to multiple locations

Field Based Services:

- Basic Deployment consists of the hardware installation, cabling and firmware update during normal standard business hours. Basic Deployment is traditionally sold to Competency Enabled Partners. Competency enabled partners often have Dell do the hardware installation while they complete the software configuration.
- ProDeploy consists of your hardware installation and configuration of the software using offshore resources. ProDeploy is great for customers who are price sensitive or who are remote from their data centers and don't require an onsite presence.
- ProDeploy Plus will give you in-region or onsite resources to complete the engagement for the customer. It also comes with additional features such as Post Deployment Configuration Assistance and Training Credits.



ProDeploy Infrastructure Suite | Factory services

		PACTORT DASED SERVICES	
		ProDeploy Factory Configuration	ProDeploy Rack Integration
and the second se	Single point of contact for project management	01	•
	RAID, BIOS and IDRAC configuration		
sset configuration	Firmware freeze		
	Asset Tagging and Reporting	0)	
	Customer system image	R	0
	Site readiness review and implementation planning	-	
and a second second second	Hardware racking and cabling	-	
actory implementation	SAM engagement for ProSupport Plus entitled accounts/devices		
	Deployment verification, documentation, and knowledge transfer	•	0-
	White glove logistics		-
	Onsite final configuration		Onsite add-on
elivery	Install support software and connect with Dell Technologies		Onsite add-on
	Basic Deployment	Optional onsite installation	
	Online collaborative environment for planning managing and tracking delivery		•

Figure 44. ProDeploy Infrastructure Suite - Factory services

ProDeploy Infrastructure Suite | Field services

		Basic Deployment	ProDeptoy	ProCepio Plas
	Single point of contact for project management.	•		In-region
	Etto readmoso review			•
e.o.coleniane	Implementation planning*	- II	•	
	SAM engagement for ProSupport Plus entitled devices	2	- 10 10	•
	Deployment service hours	Business hours	24.07	24x7
Deployment	Oneite hardware installation and packaging material remova? erremote guidance for hardware installation ¹	·	Rémote puidance or onsite	Onaite
Sector Sector	Install and configure system software		Rentatio	consta
	Install sapport software and connect with Deli Technologies			•
	Project documentation with knowledge transfer		•	
	Deployment vertication	4	•	
	Configuration data transfer to Dell Technologies technical support	e: 1 fi	•	(O)
Post-deployment	30-days of post-deployment configuration assistance			•
	Training credits for Dell Technologies Education Services			1 • 1
Online oversight	Online collaborative environment in TechDiract for planning, managing and backing delivery ²		- •	

Proceeding removal included with unsite two/ware installation included with ProDeploy or PhoDeploy Plus, Not included with Basic Deployment

Figure 45. ProDeploy Infrastructure Suite - Field services

Dell ProDeploy Plus for Infrastructure

From beginning to end, ProDeploy Plus provides the skill and scale that is must successfully perform demanding deployments in today's complex IT environments. Certified Dell experts start with extensive environmental assessments and detailed migration



planning and recommendations. Software installation includes set up of our enterprise connectivity solution (secure connect gateway) and OpenManage system management utilities.

Postdeployment configuration assistance, testing, and product orientation services are also available.

Dell ProDeploy for Infrastructure

ProDeploy provides full-service installation and configuration of both server hardware and system software by certified deployment engineers including set up of leading operating systems and hypervisors as well our enterprise connectivity solution (secure connect gateway) and OpenManage system management utilities. To prepare for the deployment, we conduct a site readiness review and implementation planning exercise. System testing, validation, and full project documentation with knowledge transfer complete the process.

Dell Basic Deployment

Basic Deployment delivers worry-free professional installation by experienced technicians who know Dell servers inside and out.

Additional Deployment Services

You can tailor the ProDeploy Infrastructure Suite offer to meet your customer's unique needs by leveraging "Additional Deployment Time." ADT will cover additional tasks above the normal scope of the standard offers. ADT can be sold for Project Management or Technical Resources and is sold as blocks of four hours remote or eight hours on-site.

Dell ProDeploy for HPC (available in US/Canada only. All other regions use custom)

HPC deployments require specialists that understand that cutting edge is yesterday's news. Dell deploys the world 's fastest systems and understands the nuances that make them perform. ProDeploy for HPC provides:

- Global team of dedicated HPC specialists
- Proven track record, thousands of successful HPC deployments
- Design validation, benchmarking, and product orientation

Learn more at Dell.com/HPC-Services.



ProDeploy Expansion for HPC

*Available as standard SKUs in US & Canada and as custom quote in APJC, EMEA, LATAM



Supplemental Deployment Services

Additional ways to expand scope or deploy for unique scenarios.

Two Host Adder (requires PD/PDP)

Deploying new storage, compute, or networking devices may require interconnection to other servers (also called hosts). The Dell delivery team will set up four hosts per device as part of every ProDeploy service. For example, if the customer is buying two storage arrays the ProDeploy service will automatically include connectivity of four hosts each (4x2=8 total hosts per project since there are two devices). This supplemental "Two Host Adder" service provides for the configuration of additional hosts above what is already provided as part of the ProDeploy service. In many cases, customers can work with us while we set up the included hosts, so they may understand how to do the rest themselves. Always ask the customer how many hosts are being connected and sell the host adder depending on the customer's technology skillset. Note that this service applies to the connectivity of Dell devices not 3rd party devices.

Additional Deployment Services (ADT) – sold with or without PD/PDP

You can expand the scope of a ProDeploy engagement leveraging Additional Deployment Time (ADT). ADT covers additional tasks above the normal deliverables of the ProDeploy offers. ADT can also be used as a standalone service without ProDeploy. SKUs are available for both Project Management and Technical Resource Expertise. SKUs are sold as blocks of four hours remote or eight hours onsite. The delivery team can help in scoping the number of hours required for additional tasks.

Data Migration Services

Migrating data sets is no easy task. Our experts use proven tools and process to streamline data migrations and avoid compromising data. A customer project manager works with our experienced team of experts to create a migration plan. Data migration is part of every technology upgrade, platform change, and shift to the cloud. You can rely on Dell data migration services to perform a seamless transition.



Residency Services

Certified technical professionals act like an extension of your IT staff to enhance internal capabilities and resources and help you realize faster adoption and maximized ROI of new technology. Residency Services help customers transition to new capabilities quickly by leveraging specific technology skill sets. Residency experts can provide post implementation management and knowledge transfer that is related to a new technology acquisition or day-to-day operational management of the IT infrastructure.

- Global experts available to serve in-person (onsite) or virtual (remote)
- Engagements starting at 2 weeks with flexibility to adjust
- Residency is available for project management needs, and many different technology skills sets such as: Server, storage, Gen AI, networking, security, multi-cloud, data mgmt., and modern workforce application residents

Unique Deployment Scenarios

Custom Deployment Services

When a deployment is beyond the scope of the ProDeploy Infrastructure Suite, you can turn to the custom deployment services team to address complex implementation scenarios and customer unique requirements. The Dell custom deployment team is staffed with solution architects who will assist with customer scoping calls to define the project and develop the statement of work. Custom services can handle a wide range of deployments that can be performed in the factory or onsite. All custom engagement services are requested through SFDC.

ProDeploy FLEX

ProDeploy Flex is a modular service and a powerful tool for you to attach more services and improve revenue and margins. The ProDeploy Flex modular offer allows sales teams to build and better tailor services by mixing factory and field delivery options. You can also select special deployment scenarios without going to the custom order desk. FLEX is ideal for unique deployments where ProDeploy or ProDeploy Plus are not an adequate answer to the customer needs. Key features of ProDeploy FLEX :

- Build deployment quotes using modular, selectable features for both hardware and software.
- The system automatically scales pricing based on volume.
- Ideal for customers who require NativeEdge Orchestrator or edge deployments.
- Ability to add deployment services to third-party networking devices.

Deployment of HPC

High-Performance Computing (HPC) implementations require specialists that understand advanced feature sets. Dell deploys the world 's fastest systems and understands the nuances that make them perform. HPC deployments are most often scoped as custom service engagements, however we can do smaller HPC clusters under 300 nodes using a standard ProDeploy SKU. Any standard SKU for HPC deployment will be sold as one base SKU per cluster (ProDeploy for HPC Base) along with one ProDeploy for HPC Add-on for each device in the cluster (server nodes and switches).

Scope of ProDeploy for HPC:

(i) NOTE: Available as standard SKUs in US and Canada. Custom Service would be required for all other regions.





- Non-Tied BASE SKU
- 1 SKU per new cluster
- (regardless of cluster size)

Figure 47. Standard deliverables of ProDeploy for HPC

HPC Add-on for Nodes

- Rack & Stack Server Nodes
- · Professionally labeled cabling
- · BIOS configured for HPC
- · OS installed
- Per node
- Tied & Non-Tied Add-on SKUs
- 1 SKU/asset
- If over 300 nodes use custom quote

Build HPC solutions for your unique requirements Choose ProDeploy for HPC or Custom deploy

ProDeploy service includes configuration of most OS, cluster mgmt., networking and benchmarking



Notes related to networking above: Omni-Path is no longer an Intel Product, but is now distributed by a company called Cornelis, and Mellanox was purchased by Nvidia, and now goes by Nvidia Networking.

Figure 48. Visual view of HPC deployment options to include hardware and software

DAY 2 – Automation Services with Ansible

Dell solutions are built as "automation ready" with integrated APIs (Application Programming Interfaces) to allow customers to programmatically call actions on the product through code. Although Dell has published Anisble automation use cases, some customers need additional assistance with GitOps. By the end of the service, the customer will have the foundational

Appendix D: Service and support 77



components required to accelerate automation and understand how the programming works together: Day 1 and Day 2 use case value automation scripts (ansible modules), CI/CD tool (Jenkins), and Version control (Git).

Dell Technologies Consulting Services

Our expert consultants help customers transform faster, and quickly achieve business outcomes for the high value workloads Dell PowerEdge systems can handle. From strategy to full-scale implementation, Dell Technologies Consulting can help determine how to perform IT, workforce, or application transformation. We use prescriptive approaches and proven methodologies that are combined with portfolio and partner ecosystem of Dell Technologies to help achieve real business outcomes. From multi cloud, applications, DevOps, and infrastructure transformations, to business resiliency, data center modernization, analytics, workforce collaboration, and user experiences-we are here to help.

Dell Managed Services

Some customers prefer Dell to manage the complexity and risk of daily IT operations, Dell Managed Services utilizes proactive, Al enabled delivery operations and modern automation to help customers realize desired business outcomes from their infrastructure investments. With these technologies, our experts run, update and fine-tune customer environments aligned with service levels, while providing environment-wide and down-to-the-device visibility. There are two types of managed service offers. First the outsourcing model or CAPEX model where Dell manages the customer owned assets using our people and tools. The second is the as-a-Service model or OPEX model called APEX. In this service, Dell owns all technology and all the management of it. Many customers will have a blend of the two management types depending on the goals of the organization.

Managed	Outsourcing or CAPEX model	APEX Service or OPEX model
We manage using our pe • Managed o • Technology • End-user (I • Service des • Cloud Man • Office365 o	your technology ople and tools. ¹ letection and response* / Infrastructure PC/desktop) sk operations aged (Pub/Private) or Microsoft Endpoint	We own all technology so you can off-load all IT decisions. • APEX Cloud Services • APEX Flex on Demand elastic capacity • APEX Data Center Utility pay-per-use model

* Managed detection and response covers the security monitoring of laptops, servers, & virtual servers. Min. 50 devices combined. No Networking or Storage-only systems [SAN/NAS]. Available in 32 countries. Details here.

Figure 49. Dell Managed Services

Managed Detection and Response (MDR)

Dell Technologies Managed Detection and Response (MDR) is powered by Secureworks Taegis XDR software platform. MDR is a managed service that secures the customer's IT environment against malicious actors and provides remediation if and when a threat is identified. When a customer purchases MDR, they will receive the following features from our team:

- Dell badge resources
- Agent rollout assistance to help deploy the Secureworks Endpoint Agent
- 24x7 threat detection & investigation
- Up to 40hrs per quarter of response and active remediation activities
- If the customer experiences a breach, we will provide up to 40hrs per year of Cyber incident response initiation
- Quarterly reviews with the customer to review the data



Dell Technologies Education Services

Build the IT skills required to influence the transformational outcomes of the business. Enable talent and empower teams with the right skills to lead and perform transformational strategy that drives competitive advantage. Leverage the training and certification required for real transformation.

Dell Technologies Education Services offers PowerEdge server training and certifications that are designed to help customers achieve more from their hardware investment. The curriculum delivers the information and the practical, firsthand skills that their team must confidently install, configure, manage, and troubleshoot Dell servers.

To learn more or register for a class today, see Education.Dell.com.

Appendix D: Service and support 79

≡			⊕ Q.
os Intel®	/ Processadores	Intel®	/ Intel® Xeon® Processors
	intel Processador Intel® Xeon® S Xeon 27.5 M 2.00 CH	ilver 4416+	
	cache de 37,5 M, 2,00 GHz		
	rocessador Intel® Xeon® S	ilver 4416+	
	cache de 37,5 M, 2,00 GHz		
	Adicionar para comparar		
	Especificações		
	Baixe as especificações $_{\rm V}$		
	Essenciais		
	Coleção de produtos	Processadores escaláveis Intel® Xeon® da 4ª Geração	
	Codinome	Produtos com denominação anterior Sapphire Rapids	
	Segmento vertical	Server	
	Número do processador ③	4416+	
	Litografia ③	Intel 7	
	Preço recomendado para o cliente 🕥	\$1176.00-\$1186.00	
	Especificações da CPU		
	Número de núcleos ③	20	
	Total de threads ③	40	
	Frequência turbo max 🧿	3.90 GHz	
	Frequência base do processador 🕥	2.00 GHz	
	Cache ③	37.5 MB	
	Velocidade do Intel® UPI	16 GT/s	
	№ de links de UPI ③	2	
	TDP ③	165 W	



Status	Launched			
Data de introdução 🎯	Q1'23			
Status de manutenção ③	Baseline Servicing			
Opções integradas disponíveis ③	Sim			
Condições de uso ③	Server/Enterprise			
Especificações de memória				
Tamanho máximo de memória (de acordo com o tipo de memória) ③	4 TB			
Tipos de memória ③	Up to DDR5 4000 MT/s 1DPC and 2DPC			
Nº máximo de canais de memória 🕥	8			
Compatibilidade com memória ECC $^{^{\dagger}}$ ③	Sim			
Opções de expansão				
Escalabilidade	2S			
Revisão de PCI Express ③	5			
№ máximo de linhas PCI Express ③	80			
Especificações de encapsulamento				
Soquetes suportados ③	FCLGA4677			
Transportadora de pacotes	E1B			
DTS Max	94 °C			
T _{CASE} ③	82			
Tamanho do pacote	77.5mm x 56.5mm			
Atualizações disponíveis do Intel® On Demand				

Activation Model Products	QAT	DLB	DSA	IAA
Communications & Storage Suite 2	2	2		
SGX512				
4				•
Tecnologias avançadas				
Ativação do recurso Intel® On Deman	d 3	Sim		
Intel® QuickAssist Technology (QAT)		1 default devices		
Intel® Dynamic Load Balancer (DLB)		1 default devices		
Intel® Data Streaming Accelerator (DS	SA)	1 default devices		
Intel® In-memory Analytics Accelerat	or (IAA)	1 default devices		
Intel® Advanced Matrix Extensions (A	MX)	Sim		

Inserido ao protocolo **22.951.206-4** por: **Pedro Henrique Golin Linhares** em: 31/03/2025 12:49. A autenticidade deste documento pode ser validada no endereço: https://www.eprotocolo.pr.gov.br/spiweb/validarDocumento com o código: fc1c44dfe0686d5547ffd691199be699.



Intel® Deep Learning Boost (Intel® DL Boost) ⑦	Sim
Intel® Resource Director Technology (Intel® RDT) ⑦	Sim
Tecnologia Intel® Speed Shift ③	Sim
Tecnologia Intel® Turbo Boost ‡ 🏵	2.0
Tecnologia Hyper-Threading Intel® ‡ 🍞	Sim
Intel® TSX-NI 💿	Sim
Intel® 64 [‡]	Sim
Extensões do conjunto de instruções 🕥	Intel® AMX, Intel® SSE4.2, Intel® AVX, Intel® AVX2, Intel® AVX-512
№ de unidades de FMA de AVX-512 🏾	2
Segurança e confiabilidade	
Intel® Software Guard Extensions (Intel®SGX) ⑦	Yes with Intel® SPS
Tamanho máximo do cache de página de enclave (EPC) para Intel® SGX	64 GB
Intel® Crypto Acceleration ③	Sim
Aceleração de software Intel® QuickAssist	Sim
Suporte para Resiliência de firmware de plataforma Intel®	Sim
Intel® Control-Flow Enforcement Technology ⑦	Sim
Intel® Total Memory Encryption ③	Sim
Novas instruções Intel® AES 🕥	Sim
Intel® OS Guard	Sim
Intel * Trusted Execution Technology $^{^{\dagger}}$ ()	Sim
Bit de desativação de execução ‡ 🕉	Sim
Intel® Boot Guard ③	Sim
Controle de Execução baseado em Modo (MBEC — Mode-based Execute Control) ③	Sim
Tecnologia de virtualização Intel® (VT-x) † 🍞	Sim
Tecnologia de virtualização Intel® para E/S dirigida (VT-d) † ③	Sim
Intel® VT-x com Tabelas de páginas estendidas (EPT) [†] ⑦	Sim

Todas as informações fornecidas estão sujeitas a alterações a qualquer momento, sem aviso prévio. A Intel pode alterar o ciclo de vida da fabricação, as especificações e as descrições dos produtos a qualquer momento, sem aviso prévio. As informações aqui contidas são fornecidas "no estado em que se encontram" e a Intel não atribui qualquer declaração ou garantias relacionadas à precisão das informações, nem sobre os recursos dos produtos, disponibilidade, funcionalidade ou compatibilidade dos produtos listados. Para obter mais informações sobre os produtos ou sistemas, entre em contato com o fornecedor do sistema.

Inserido ao protocolo **22.951.206-4** por: **Pedro Henrique Golin Linhares** em: 31/03/2025 12:49. A autenticidade deste documento pode ser validada no endereço: https://www.eprotocolo.pr.gov.br/spiweb/validarDocumento com o código: fc1c44dfe0686d5547ffd691199be699.



Intel classifications are for general, educational and planning purposes only and consist of Export Control Classification Numbers (ECCN) and Harmonized Tariff Schedule (HTS) numbers. Any use made of Intel classifications are without recourse to Intel and shall not be construed as a representation or warranty regarding the proper ECCN or HTS. Your company as an importer and/or exporter is responsible for determining the correct classification of your transaction.

Consulte a Ficha técnica para obter definições formais de propriedades e recursos de produtos.

‡ Este recurso pode não estar disponível em todos os sistemas de computação. Verifique com o fornecedor do sistema para determinar se seu sistema oferece este recurso ou consulte as especificações de seu sistema (motherboard, processador, chipset, alimentação, HDD, controle gráfico, memória, BIOS, drivers, monitor de máquina virtual [VMM], software de plataforma e/ou sistema operacional) para saber sobre a compatibilidade do recurso. A funcionalidade, o desempenho e outros benefícios deste recurso podem variar, dependendo das configurações do sistema.

Os números dos processadores Intel não são indicação de desempenho. Os números dos processadores diferenciam recursos dentro de cada família de processador, e não entre famílias diferentes de processadores. Consulte https://www.intel.com.br/content/www/br/pt/processors/processor-numbers.html

para obter mais detalhes.

O RCP (Recommended Customer Price, preço recomendado para o cliente) é o guia de preços somente para produtos Intel. Os preços são para clientes diretos da Intel, representam geralmente as quantidades de compra de 1.000 unidades, e estão sujeitos a alterações sem aviso prévio. Os preços podem variar para outros tipos de pacotes e quantidades de envio. Na venda por atacado, o preço corresponde à unidade. Listar os índices RCP não constituí uma oferta oficial da Intel.

O TDP máximo e do sistema se baseiam nos piores casos. O TDP real pode ser inferior, se nem todas as E/Ss para chipsets forem utilizadas.

SKUs "anunciados" ainda não estão disponíveis. Favor consultar a data de lançamento para a disponibilidade no mercado.

Frequência máxima de turbo refere-se à frequência máxima do processador de núcleo único que pode ser atingida com a Tecnologia Intel® Turbo Boost. Mais informações estão disponíveis no site https://www.intel.com/content/www/br/pt/architecture-and-technology/turbo-boost/turbo-boost-technology.html

Consulte https://www.intel.com.br/content/www/br/pt/architecture-and-technology/hyper-threading/hyper-threading-technology.html? wapkw=hyper+threading

para obter mais informações, incluindo detalhes sobre quais processadores são compatíveis com a Tecnologia Hyper-Threading Intel[®].

Os processadores compatíveis com a computação de 64 bits na arquitetura Intel® requerem BIOS habilitados para arquitetura Intel 64.

Alguns produtos suportam as novas instruções AES com uma atualização da Configuração do processador, em particular, i7-2630QM/i7-2635QM, i7-2670QM/i7-2675QM, i5-2430M/i5-2435M, i5-2410M/i5-2415M. Favor entrar em contato com o OEM para o BIOS que inclui a mais recente atualização da Configuração do processador.

Informações sobre a empresa

Nosso compromisso

Inclusão

Relações com investidores

Fale conosco

Sala de imprensa

Mapa do site

Empregos

f

in

© Intel Corporation Termos de uso *Marcas comerciais Cookies Privacidade Transparência da cadeia de fornecimento

As tecnologias Intel[®] podem exigir ativação de hardware, software específico ou de serviços. // Nenhum produto ou componente pode ser totalmente seguro. // Os seus custos e resultados podem variar. // O desempenho varia de acordo com o uso, a configuração e outros fatores. // Veja nossos Avisos e isenções de responsabilidade legais completos

 \mathbb{X}

. // A Intel está comprometida em respeitar os direitos humanos e evitar cumplicidade com abusos de direitos humanos. Consulte Princípios Globais de Direitos Humanos da Intel. Os produtos e software da Intel são destinados a serem utilizados apenas em aplicações que não causem ou

contribuam com a violação de um direito humano reconhecido internacionalmente.

r para o conteúdo principal



intel.

Inserido ao protocolo **22.951.206-4** por: **Pedro Henrique Golin Linhares** em: 31/03/2025 12:49. A autenticidade deste documento pode ser validada no endereço: https://www.eprotocolo.pr.gov.br/spiweb/validarDocumento com o código: fc1c44dfe0686d5547ffd691199be699.