

# PROPOSTA TÉCNICA E COMERCIAL PREGÃO ELETRÔNICO RP Nº 40/2019



## TRIBUNAL REGIONAL FEDERAL DA PRIMEIRA REGIÃO – TRF 1ª REGIÃO

LTA-RH Informática Comércio, Representações Ltda.  
ST SCN QUADRA 04 - BLOCO B, SALA 1251/1252  
ASA NORTE – Brasília/DF, CEP 70714-900  
CGC-MF n.º 94.316.916/0005-22  
Inscrição Estadual n.º 07.572.077/002-71  
Fone/Fax: 51-3382.7700 / 51-3382.7722  
E-mail: [comercial@lta-rh.com.br](mailto:comercial@lta-rh.com.br)



[www.lta-rh.com.br](http://www.lta-rh.com.br) | [comercial@lta-rh.com.br](mailto:comercial@lta-rh.com.br)

**Matriz** | Av. Ipiranga, 2640 | Santa Cecília | Porto Alegre | RS | Brasil | CEP 90610-000 | (51) 3382.7700/3094.1500  
**Filial DF** | SCN, Quadra 04 | Bloco B | Sala 1251/1252 | Distrito Federal | DF | Brasil | CEP:70714-900 | (61) 3034-3004  
**Filial SP** | Av. Paulista, 1636 | Conj.706 | Bela Vista | São Paulo | SP | Brasil | CEP 01310-200 | (11) 2391-9461  
**Filial RJ** | Praia de Botafogo 501 | Blc I Sala 101 | Botafogo | Rio de Janeiro | RJ | CEP 22.250-040 | (21) 2586-6000  
**Filial PR** | Rua Comendador Araújo 499 | CONJ 1007 | Centro | Curitiba | PR | CEP: 80.420-000 | (41) 99104-3240  
**Filial MG** | Av. Do Contorno, 6594 | 705 | Belo Horizonte | MG | CEP 30110-044 | (31) 3555-3477

Brasília, 1º de novembro de 2019.

Ao  
TRIBUNAL REGIONAL FEDERAL DA PRIMEIRA REGIÃO – TRF 1ª REGIÃO  
Ref.: PREGÃO ELETRÔNICO RP Nº 40/2019  
Proposta N.º 281/19

Prezado Senhor:

Apresentamos a V.Sa. a nossa “*proposta técnica e comercial*” para fornecimento de material objeto da presente licitação.

Cumpre-nos informar-lhes que examinamos cuidadosamente os documentos da licitação, inteirando-nos dos mesmos, para elaboração da presente proposta.

Portanto, na expectativa de decisão favorável, colocamo-nos ao inteiro dispor de V.s.as. para quaisquer esclarecimentos.

Atenciosamente,

ALEXANDER C. BARCELOS  
DIRETOR COMERCIAL



[www.lta-rh.com.br](http://www.lta-rh.com.br) | [comercial@lta-rh.com.br](mailto:comercial@lta-rh.com.br)

**Matriz** | Av. Ipiranga, 2640 | Santa Cecília | Porto Alegre | RS | Brasil | CEP 90610-000 | (51) 3382.7700/3094.1500

**Filial DF** | SCN, Quadra 04 | Bloco B | Sala 1251/1252 | Distrito Federal | DF | Brasil | CEP:70714-900 | (61) 3034-3004

**Filial SP** | Av. Paulista, 1636 | Conj.706 | Bela Vista | São Paulo | SP | Brasil | CEP 01310-200 | (11) 2391-9461

**Filial RJ** | Praia de Botafogo 501 | Blc I Sala 101 | Botafogo | Rio de Janeiro | RJ | CEP 22.250-040 | (21) 2586-6000

**Filial PR** | Rua Comendador Araújo 499 | CONJ 1007 | Centro | Curitiba | PR | CEP: 80.420-000 | (41) 99104-3240

**Filial MG** | Av. Do Contorno, 6594 | 705 | Belo Horizonte | MG | CEP 30110-044 | (31) 3555-3477

**ANEXO I**  
**SÍNTESE DO TERMO DE REFERÊNCIA**

**Ao**  
**TRIBUNAL REGIONAL FEDERAL DA PRIMEIRA REGIÃO – TRF 1ª REGIÃO**  
**Ref.: PREGÃO ELETRÔNICO RP Nº 40/2019**  
**Proposta N.º 281/19**

**1. DO OBJETO**

1.1. O presente Termo de Referência tem por objeto a contratação de empresa especializada no fornecimento de equipamentos Servidores tipo Rack e Solução de Hiperconvergência, com assistência técnica pelo período de 60 (sessenta) meses e serviço de instalação, para atender às necessidades das subseções e seções judiciárias, bem como do Tribunal Regional Federal da 1ª Região - TRF1 de acordo com condições, especificações e quantidades constantes deste Termo e Anexo I-A.

**3. DOS ITENS E QUANTITATIVOS**

3.1. Segue abaixo, a planilha contendo os quantitativos a serem registrados:

LOTE	ITEM	DESCRIÇÃO	MEDIDA	QTD	CÓD. SIASG
1	01	Solução de hiperconvergência	Unidade	82	BR459964
	02	Instalação - Solução de hiperconvergência	Serviço	82	BR027111

**4. DA VISTORIA**

4.1. A LICITANTE será possibilitado vistoriar os locais de prestação dos serviços, no período das 10h às 18h, sob agendamento pelo telefone (61) 3314- 1695 ou pelo e-mail ditech@trf1.jus.br.

4.1.1. A faculdade da vistoria visa a realização de visita às dependências do Tribunal Regional Federal da 1ª Região, mediante agendamento, para esclarecimentos de dúvidas relacionadas ao objeto licitado, bem como para verificar os locais de instalação dos equipamentos, devendo em todo caso apresentar Declaração de Vistoria Técnica ou renúncia ao direito de vistoria, conforme modelo constante no Anexo IV deste Edital. A vistoria tem por finalidade avaliar as condições das instalações atualmente existentes e tornar registrado o pleno conhecimento das proponentes acerca das dificuldades para a execução do objeto e avaliar a forma de melhor executar o objeto, comparativamente aos recursos técnicos de que dispõe.

4.2. A DITEC/SECIN intermediará o agendamento da vistoria nas localidades indicadas pela LICITANTE, permanecendo à disposição, por telefone e no horário agendado para prestar quaisquer informações complementares que a LICITANTE julgue necessário, estando a relação de contatos das localidades consolidada no Item 05 deste Anexo.

4.3. Caso a LICITANTE opte por não realizar a vistoria, deverá declarar que tem pleno conhecimento dos ambientes tecnológicos do CONTRATANTE em que serão entregues os equipamentos e prestados os serviços, assumindo total responsabilidade pela ocorrência de eventuais prejuízos, em virtude de sua omissão na verificação dos ambientes tecnológicos do CONTRATANTE, com vistas a proteger o interesse da Administração na fase de execução do contrato, nos termos do Acórdão 1.174/2008 – Plenário - TCU.

4.4. Em havendo a realização da vistoria A LICITANTE deverá apresentar juntamente com a proposta, a Declaração de Vistoria Técnica emitida pela DITEC/SECIN.

**5. DO LOCAL DA ENTREGA E PRESTAÇÃO DOS SERVIÇOS**

5.1. Endereço para entrega dos equipamentos e execução dos serviços conforme segue:



[www.lta-rh.com.br](http://www.lta-rh.com.br) | [comercial@lta-rh.com.br](mailto:comercial@lta-rh.com.br)

**Matriz** | Av. Ipiranga, 2640 | Santa Cecília | Porto Alegre | RS | Brasil | CEP 90610-000 | (51) 3382.7700/3094.1500

**Filial DF** | SCN, Quadra 04 | Bloco B | Sala 1251/1252 | Distrito Federal | DF | Brasil | CEP:70714-900 | (61) 3034-3004

**Filial SP** | Av. Paulista, 1636 | Conj.706 | Bela Vista | São Paulo | SP | Brasil | CEP 01310-200 | (11) 2391-9461

**Filial RJ** | Praia de Botafogo 501 | Blc I Sala 101 | Botafogo | Rio de Janeiro | RJ | CEP 22.250-040 | (21) 2586-6000

**Filial PR** | Rua Comendador Araújo 499 | CONJ 1007 | Centro | Curitiba | PR | CEP: 80.420-000 | (41) 99104-3240

**Filial MG** | Av. Do Contorno, 6594 | 705 | Belo Horizonte | MG | CEP 30110-044 | (31) 3555-3477

LOCALIDADE		ENDEREÇO	FONE
DF	TRF1 - Tribunal Regional Federal da 1ª Região	SAL/SUL Quadra 2, Bloco K, Edifício Sede II, Praça dos Tribunais Superiores - Brasília - DF - CEP: 70.070-900	(61) 3314-5225
AC	SJAC - Seção do Rio Branco	Cidade da Justiça, Rod BR 307, Km 9, n. 4.090 Bairro Boca da Alemanha - Cruzeiro do Sul - AC - CEP: 69.980-000	(68) 3214-2000
	Subseção de Cruzeiro do Sul	Cidade da Justiça, Rod BR 307, Km 9, n. 4.090 Bairro - Boca da Alemanha - Cruzeiro do Sul - AC - CEP: 69.980-000	(68)3311-1770
AM	SJAM - Seção de Manaus	Av. André Araújo, n. 25 - Bairro Aleixo - Manaus - AM - CEP: 69.060-000.	(92) 3612-3300
	Subseção de Tabatinga	Avenida Aires da Cunha, n. 48 Bairro Ibirapuera Tabatinga - AM - CEP: 69.640-000	(97) 3412-5205
	Subseção de Tefé	Rua Daniel Sevalho s/n - Centro - Tefé - AM - CEP: 69.470-000	(97) 3343-2870
AP	SJAP - Seção de Macapá	Rodovia Norte-Sul, s/n, Infraero II - Macapá - AP - CEP: 68.908-911.	(96) 3214-1500
	Subseção de Laranjal do Jari	Rua Rio Jari, n. 1112 - Bairro do Agreste AP. Laranjal do Jari - AP - CEP 68.920-000	(96) 3621-1534
	Subseção de Oiapoque	Rua Santos Dumont, n. 769 Centro. - Oiapoque - AP CEP: 68.980-000	(96) 3521-1618
BA	SJBA - Seção de Salvador	Fórum Teixeira de Freitas, Av. Ulysses Guimarães, n. 2799 - Bairro Sussuarana - Centro Administrativo da Bahia - Salvador-BA - CEP: 41.213-000.	(71) 3617-2600
	Subseção de Alagoinhas	Endereço: R. Silva Jardim, s/n (prédio do INSS), 2º andar - Parque São Jorge. - Alagoinhas - BA - CEP: 48.060-000	(75) 3422-6729
	Subseção de Barreiras	Rua Aníbal Barbosa, s/n, Bairro Centro, 5º andar - Fórum Tarcílio Vieira de Melo. - Barreiras - BA - CEP: 47.800-163	(77) 3611-8391
	Subseção de Bom Jesus da Lapa	Av. Agenor Magalhães, s/n Bairro Mirante da Lapa Bom Jesus da Lapa - BA - CEP 47.600-000	(77) 3481-2026
	Subseção de Campo Formoso	Praça da Bandeira nº 95, Centro - Campo Formoso - BA - CEP: 44.790-000	(74) 3645-1987
	Subseção de Eunápolis	Av. Antonio Carlos Magalhães, n. 160, Bairro Dinah Borges Eunápolis - BA - CEP: 45.820-540	(73) 3261-7070
	Subseção de Feira de Santana	Rua Turquia, s/n - Ponto Central - Feira de Santana - Ba CEP: 44.075.295	(75) 3221-6274
	Subseção de Guanambi	Av. Santos Dumont, n. 325 - Centro - Guanambi - BA - CEP: 46.430-000	(77) 3451-3788
	Subseção de Ilhéus	Rua Ministro José Cândido, n. 80 - Centro. Ilhéus - BA - CEP: 45.653-542	(73) 3634-2950
	Subseção de Irecê	Avenida Sol Poente, s/n, bairro Asa Norte, Novo Forum da Comarca de Irecê, 4º andar. Irecê - BA - CEP 44900-000	(74) 3641-3910
	Subseção de	Av. Amélia Amado, n. 331, Centro (antigo CNPC)	(73) 3215-3388



[www.lta-rh.com.br](http://www.lta-rh.com.br) | [comercial@lta-rh.com.br](mailto:comercial@lta-rh.com.br)

**Matriz** | Av. Ipiranga, 2640 | Santa Cecília | Porto Alegre | RS | Brasil | CEP 90610-000 | (51) 3382.7700/3094.1500

**Filial DF** | SCN, Quadra 04 | Bloco B | Sala 1251/1252 | Distrito Federal | DF | Brasil | CEP:70714-900 | (61) 3034-3004

**Filial SP** | Av. Paulista, 1636 | Conj.706 | Bela Vista | São Paulo | SP | Brasil | CEP 01310-200 | (11) 2391-9461

**Filial RJ** | Praia de Botafogo 501 | Blc I Sala 101 | Botafogo | Rio de Janeiro | RJ | CEP 22.250-040 | (21) 2586-6000

**Filial PR** | Rua Comendador Araújo 499 | CONJ 1007 | Centro | Curitiba | PR | CEP: 80.420-000 | (41) 99104-3240

**Filial MG** | Av. Do Contorno, 6594 | 705 | Belo Horizonte | MG | CEP 30110-044 | (31) 3555-3477

	Itabuna	- Itabuna - BA - CEP: 45.600-033	
	Subseção de Jequié	Rua Gildélito Ferraz, s/, Jequezinho. Jequié -BA - CEP: 45.208-415	(73) 3525-6355
	Subseção de Juazeiro	Rua Dom Pedro I s/n, Bairro João XXIII - Juazeiro - BA - CEP 48.900-230	(74) 3611-7970
	Subseção de Paulo Afonso	Rua do Gangorra, Quadra 12, n. 148-A - Alves de Souza Paulo Afonso - BA - CEP 48.608-240	(75)3281-2387
	Subseção de Teixeira de Freitas	Av. Marechal Castelo Branco, n. 229, Centro - Teixeira de Freitas - BA - CEP: 45985-160	(73) 3291-1731
	Subseção de Vitória da Conquista	Rua Ministro Hermes Lima, S/N - Cidade Universitária. Vitória da Conquista - BA. CEP: 45.029-260.	(77) 3423-8900
DF	SJDF - Seção do Distrito Federal	Edifício-Sede I SAS - Praça dos Tribunais Superiores Quadra 2, bloco G, Lote 8 - Brasília - DF - CEP: 70.070-933	(61) 3221-6310
GO	SJGO - Seção de Goiânia	Rua 19 n. 244, Centro. Goiânia - GO - CEP: 74.030-090	(62) 3226-1515
	Subseção de Anápolis	Av. Universitária, Quadra 02, Lote 05, Jardim - Bandeirantes. Anápolis - GO - CEP: 75.083-035	(62) 4015 - 8609
	Subseção de Aparecida de Goiânia	Av. J-2 com Rua J-17 e Rua J-4, Quadra 35, Lote 1/4, Mansões Paraíso - Aparecida de Goiânia - GO - CEP: 74.952-180	(62) 3625 - 8009
	Subseção de Formosa	Endereço: Rua Itiquira, n. 1000, com Rua Lindolfo Gonçalves, Setor Nordeste. Formosa - GO - CEP: 73.807-145	(61) 3631-4416
	Subseção de Itumbiara	Avenida João Paulo II, n. 185 (Térreo do Edifício do Fórum), Bairro Ernestina Borges de Andrade. Itumbiara - GO - CEP:75.528-370	(61) 2103-6410
	Subseção de Jataí	Rua Nicolau Zaidem, Qd. 45, Vila Fátima, Centro. Jataí - GO - CEP: 75.803-055	(64) 2102-2114
	Subseção de Luziânia	Rua Dr. João Teixeira, n. 596 Quadra 73, Lote 21-A, Edifício Iaci Amaral, Centro. Luziânia - GO - CEP: 72.800-440	(61) 2104-3506
	Subseção de Rio Verde	Av. Universitária, n. 359 Quadra 11, Lote 17-B, Jardim Presidente. Rio Verde - GO - CEP: 75.901-970	(64) 3211-8601
	Subseção de Uruaçu	Av. Tocantins, n. 17, Quadra 07, Lote 16 Centro - Uruaçu - GO CEP: 76.400-000	(62) 3357-1070
MA	SJMA - Seção de São Luiz	Av. Senador Vitorino Freire, n. 300Areinha - São Luís - MA - CEP 65.031-900	(98) 3214-5701
	Subseção de Bacabal	Rua Frederico Leda, n. 1910 Centro - Bacabal - MA. - CEP: 65.700-000	(99) 3627-6700
	Subseção de Balsas	Rodoviária BR 230, s/n.º, Setor Industrial, Balsas - MA - CEP: 65.800-000	(99) 3542-5561
	Subseção de Caxias	Rua 07-A, Cidade Judiciária Bairro Campo de Belém. Caxias- MA - CEP: 65.609-900	(99) 3422-0200



[www.lta-rh.com.br](http://www.lta-rh.com.br) | [comercial@lta-rh.com.br](mailto:comercial@lta-rh.com.br)

**Matriz** | Av. Ipiranga, 2640 | Santa Cecília | Porto Alegre | RS | Brasil | CEP 90610-000 | (51) 3382.7700/3094.1500

**Filial DF** | SCN, Quadra 04 | Bloco B | Sala 1251/1252 | Distrito Federal | DF | Brasil | CEP:70714-900 | (61) 3034-3004

**Filial SP** | Av. Paulista, 1636 | Conj.706 | Bela Vista | São Paulo | SP | Brasil | CEP 01310-200 | (11) 2391-9461

**Filial RJ** | Praia de Botafogo 501 | Blc I Sala 101 | Botafogo | Rio de Janeiro | RJ | CEP 22.250-040 | (21) 2586-6000

**Filial PR** | Rua Comendador Araújo 499 | CONJ 1007 | Centro | Curitiba | PR | CEP: 80.420-000 | (41) 99104-3240

**Filial MG** | Av. Do Contorno, 6594 | 705 | Belo Horizonte | MG | CEP 30110-044 | (31) 3555-3477



	Subseção de Imperatriz	Av. Tapajós, S/N - Bairro: Parque das Nações - Imperatriz-MA - CEP 65.912-900	(99) 3529-0566
MG	SJMG - Seção de Belo Horizonte	Av. Álvares Cabral, n. 1805 Bairro Santo Agostinho Belo Horizonte - MG - CEP: 30.170-001	(31) 3501-1800
	Subseção de Contagem	Avenida José Faria da Rocha, 5.021 - Bairro Eldorado. Contagem - MG - CEP: 32.310-210	(31) 3268-6300
	Subseção de Divinópolis	Praça Dom Cristiano, 298 Centro. Divinópolis - MG - CEP: 35.500-004	(37) 2101-8000
	Subseção de Governador Valadares	Rua Bárbara Heliodora, 862 - Bairro Centro. Governador Valadares - MG - CEP: 35.010-040	(33) 2101-8100
	Subseção de Ipatinga	Avenida Selim José de Sales, n. 895 - Canaã - Ipatinga - MG - CEP: 35.164-213	(31) 2109-8107
	Subseção de Ituiutaba	Rua Vinte e Oito, 1155 - Bairro Centro - Ituiutaba - MG - CEP: 38.300-032	(34) 3271-1900
	Subseção de Janaúba	Rua São João da Ponte, 1332 - Bairro São Gonçalo - Janaúba - MG - CEP: 39.440-000	(38) 3829-4950
	Subseção de Juiz de Fora	Rua Leopoldo Schmidt, n. 145 - Centro - Juiz de Fora - MG - CEP: 36.060-040	(32) 3311-1500
	Subseção de Lavras	Rua Kennedy dos Santos, 40 Bairro Santa Efigênia. - Lavras - MG - CEP: 37.200-000	(35) 3826-8702
	Subseção de Manhuaçu	Rua Duarte Peixoto, 70, Bairro Coqueiro - Manhuaçu - MG - CEP: 36.900-000	(33) 3339-2656
	Subseção de Montes Claros	Av. Deputado Esteves Rodrigues, n. 852 -Bairro Centro - Montes Claros - MG - CEP: 39.400-215	(38) 2101-8200
	Subseção de Muriaé	Rua Dr. Mário Inácio Carneiro, 535 - Centro. - Muriaé - MG - CEP: 36.889-007	(32) 3511-0100
	Subseção de Paracatu	Av. Olegário Maciel, 138 - Paracatu - MG - CEP: 38.600-000	(38) 3311-1341
	Subseção de Passos	Av. Arlindo Figueiredo n. 128 - Bairro Centro - Passos - MG - CEP: 37.902-026	(35) 3211-1160
	Subseção de Patos de Minas	Rua José de Santana, 506 - Sala 106 - Centro - Patos de Minas - MG - CEP: 38.700-052.	(34) 3818-5400
	Subseção de Poços de Caldas	Av. João Pinheiro, 1071 - Bairro Campo da Mogiana - Poços de Caldas - MG - CEP: 37.701-880.	(35) 3697-4450
	Subseção de Ponte Nova	Av. Caetano Marinho, 268 - Centro - Ponte Nova - MG - CEP: 35.430-001	(31) 3604-1002
	Subseção de Pouso Alegre	Rua Santo Antonio nº 105 Centro - Pouso Alegre - MG CEP 37.550-000	(35) 2102-1060
	Subseção de São João Del- Rei	Rua Professora Margarida Moreira Neves, n. 170 Bairro São Judas Tadeu - São João Del-Rei - MG - CEP: 36.307-066	(32) 3322-1203
Subseção de São Sebastião do Paraíso	Av. Oliveira Rezende, 662 - Brás - São Sebastião do Paraíso - MG - CEP: 37.950-000.	(35) 3411-1150	
	Subseção de Sete	Rua Santos Dumont, 140- Bairro Canaã - Sete	(31) 2106-8000



[www.lta-rh.com.br](http://www.lta-rh.com.br) | [comercial@lta-rh.com.br](mailto:comercial@lta-rh.com.br)

**Matriz** | Av. Ipiranga, 2640 | Santa Cecília | Porto Alegre | RS | Brasil | CEP 90610-000 | (51) 3382.7700/3094.1500

**Filial DF** | SCN, Quadra 04 | Bloco B | Sala 1251/1252 | Distrito Federal | DF | Brasil | CEP:70714-900 | (61) 3034-3004

**Filial SP** | Av. Paulista, 1636 | Conj.706 | Bela Vista | São Paulo | SP | Brasil | CEP 01310-200 | (11) 2391-9461

**Filial RJ** | Praia de Botafogo 501 | Blc I Sala 101 | Botafogo | Rio de Janeiro | RJ | CEP 22.250-040 | (21) 2586-6000

**Filial PR** | Rua Comendador Araújo 499 | CONJ 1007 | Centro | Curitiba | PR | CEP: 80.420-000 | (41) 99104-3240

**Filial MG** | Av. Do Contorno, 6594 | 705 | Belo Horizonte | MG | CEP 30110-044 | (31) 3555-3477

	Lagoas	Lagoas – MG. CEP: 35.700-284.	
	Subseção de Teófilo Otoni	Rua Dr. Reinaldo, Nº 105 - Centro - Teófilo Otoni – MG. - CEP: 39.800-018	(33) 3087-0109
	Subseção de Uberaba	Av. Maria Carmelita de Castro Cunha, 30 – Bairro Vila Olímpica. Uberaba – MG - CEP: 38.065-320	(34) 2103-5100
	Subseção de Uberlândia	Av. Cesário Alvim, n. 3390 - Bairro Brasil - Uberlândia – MG CEP 38.400-696	(34) 2101-3801
	Subseção de Unaí	Rua João Pinheiro, 548 - Centro - Unaí – MG. CEP: 38.610- 000.	(38)2102-1891
	Subseção de Varginha	Av. Princesa do Sul, 1855 – Rezende - Varginha – MG - CEP: 37.062-447	(35) 2105-8101
	Subseção de Viçosa	Avenida Joaquim Lopes de Farias, 505 - Santo Antônio. Viçosa – MG - CEP: 36.570-000	(31) 3611-7404
MT	SJMT - Seção de Cuiabá	Av. Rubens de Mendonça, n. 4.888, Fórum Federal JJ Rabelo, Edifício Desembargador Federal Mário Mendes – Centro Político Administrativo. - Cuiabá – MT - CEP: 78.049-942	(65) 3614-5700
	Subseção de Barra do Garças	Av. Senador Valdon Varjão 3494 Setor Industrial. - Barra do Garças – MT - CEP: 78.600-000	(66) 3402-0000
	Subseção de Cáceres	Rua Generoso Marques Leite 300 – COC - Cáceres-MT - CEP: 78.200-000	(65) 3211-6109
	Subseção de Diamantino	Rua Rui Barbosa, Quadra 30, Lote 39 São Sebastião- Diamantino - MT - CEP: 78.400-000	(65) 3336-6800
	Subseção de Juína	Av. Gabriel Muller n. 794-N, Bairro Módulo 1 - Juína – MT - CEP: 78.320-000	(66) 3524-0100
	Subseção de Rondonópolis	Av. Goiânia, 281, Bairro Santa Marta - Rondonópolis – MT - CEP 78.710-450	(66) 3902-2278
	Subseção de Sinop	Av. Julio Campos, 1230. Centro - Sinop – MT. - CEP: 78.550- 286.	(66) 3901-1250
PA	SJPA - Seção de Belém	Rua Domingos Marreiros, 598, I Bairro Umarizal - Belém – PA - CEP: 66.055-210.	(91) 3299-6100
	Subseção de Altamira	AV. Tancredo Neves, n. 100 Bairro Premem - Altamira - PA - CEP: 68.372-060	(93) 3515-2920
	Subseção de Castanhal	Rua Quintino Bocaiúva, 2.363 Centro - Castanhal – PA CEP: 68.743-010	(91) 3412-2750
	Subseção de Itaituba	Av. Paes de Carvalho, s/n Centro (mesma rua do Fórum Estadual) - Itaituba-PA - CEP 68.180-060	(93) 3518-0760
	Subseção de Marabá	Travessa Ubá, s/n, Bairro Amapá - Marabá – PA. CEP 68.502- 008	(94) 3324-2486
	Subseção de Paragominas	Av. Portugal, 3 - QD 03 - Bloco 05 - Módulo II – Paragominas – PA - CEP 68.626-080	(91) 3729-9400
	Subseção de Redenção	Av. Independência, lote n. 11, quadra 34 - Bairro Centro. Redenção – PA - CEP 68.553-055	(94) 3424-1105
	Subseção de Santarém	Avenida Barão do Rio Branco, 1893 Bairro Jardim Santarém - Santarém – PA - CEP: 68.005-396	(93) 2101-9450
	Subseção de Tucuruí	Rua 01, n. 51, Bairro Jardim Marilucy, 2º Piso - Tucuruí – PA. CEP: 68.459-490.	(94) 3787-6004



[www.lta-rh.com.br](http://www.lta-rh.com.br) | [comercial@lta-rh.com.br](mailto:comercial@lta-rh.com.br)

**Matriz** | Av. Ipiranga, 2640 | Santa Cecília | Porto Alegre | RS | Brasil | CEP 90610-000 | (51) 3382.7700/3094.1500

**Filial DF** | SCN, Quadra 04 | Bloco B | Sala 1251/1252 | Distrito Federal | DF | Brasil | CEP:70714-900 | (61) 3034-3004

**Filial SP** | Av. Paulista, 1636 | Conj.706 | Bela Vista | São Paulo | SP | Brasil | CEP 01310-200 | (11) 2391-9461

**Filial RJ** | Praia de Botafogo 501 | Blc I Sala 101 | Botafogo | Rio de Janeiro | RJ | CEP 22.250-040 | (21) 2586-6000

**Filial PR** | Rua Comendador Araújo 499 | CONJ 1007 | Centro | Curitiba | PR | CEP: 80.420-000 | (41) 99104-3240

**Filial MG** | Av. Do Contorno, 6594 | 705 | Belo Horizonte | MG | CEP 30110-044 | (31) 3555-3477

PI	SJPI - Seção de Teresina	Avenida Miguel Rosa, N. 7.315 Bairro Redenção Teresina – PI - CEP: 64.018-550.	(86) 2107-2800
	Subseção de Corrente	BR-135 - KM -49 - Zona Urbana do Município de Corrente – PI CEP: 64.980-000	(89) 3573-4000
	Subseção de Floriano	Rua Fernando Drummond n. 881 – Centro - Floriano-PI - CEP: 64.800-072	(89) 3515-6600
	Subseção de Parnaíba	Av. Humberto de Campos, 634 - Centro - Parnaíba-PI - CEP: 64.200-380	(86) 3322-1950
	Subseção de Picos	Rua Santo Antonio n. 74 – Centro - Picos-PI - CEP – 64.600- 004	(89) 2101-2800
	Subseção de São Raimundo Nonato	Rua Frade Macedo, n. 1054- Aldeia - São Raimundo Nonato – PI - CEP: 64.600 -000	(89) 3582-9600
RO	SJPA - Seção de Porto Velho	Av. Presidente Dutra n. 2203, Centro Porto Velho – RO - CEP: 76.805-902.	(69) 3211-2410
	Subseção de Guajará Mirim	Av. Duque de Caxias, 2409, Bairro Santa Luzia - Guajará Mirim – RO. - CEP: 76.850-000	(69) 3541-6348
	Subseção de Ji-Paraná	Av. Marechal Rondon, n. 935 – Centro - Ji-Paraná – RO - CEP: 76.900-081	(69) 3903-1000
	Subseção de Vilhena	Av. Brigadeiro Eduardo Gomes, 1196 - Bairro Jardim Eldorado Vilhena-RO - CEP: 76.980-000	(69) 3321-2090
TO	SJTO - Seção de Palmas	201 Norte, Conjunto 1, Lote 02A, Caixa Postal 161 – Palmas – TO - CEP: 77.001-128	(63) 3218-3809
	Subseção de Araguaina	Avenida José de Brito, Lote 05, Quadra M-12, Setor Anhanguera - Araguaina-TO - CEP: 77.818-530	(63) 2112-8202
	Subseção de Gurupi	Avenida São Paulo, esquina com Rua 10, n. 1.680 Gurupi – TO - CEP: 77.403-040.	(63) 3301-3805



[www.lta-rh.com.br](http://www.lta-rh.com.br) | [comercial@lta-rh.com.br](mailto:comercial@lta-rh.com.br)

**Matriz** | Av. Ipiranga, 2640 | Santa Cecília | Porto Alegre | RS | Brasil | CEP 90610-000 | (51) 3382.7700/3094.1500

**Filial DF** | SCN, Quadra 04 | Bloco B | Sala 1251/1252 | Distrito Federal | DF | Brasil | CEP:70714-900 | (61) 3034-3004

**Filial SP** | Av. Paulista, 1636 | Conj.706 | Bela Vista | São Paulo | SP | Brasil | CEP 01310-200 | (11) 2391-9461

**Filial RJ** | Praia de Botafogo 501 | Blc I Sala 101 | Botafogo | Rio de Janeiro | RJ | CEP 22.250-040 | (21) 2586-6000

**Filial PR** | Rua Comendador Araújo 499 | CONJ 1007 | Centro | Curitiba | PR | CEP: 80.420-000 | (41) 99104-3240

**Filial MG** | Av. Do Contorno, 6594 | 705 | Belo Horizonte | MG | CEP 30110-044 | (31) 3555-3477



## ANEXO I-A - DAS ESPECIFICAÇÕES TÉCNICAS

### 1 ITEM 1 - SOLUÇÃO DE HIPERCONVERGÊNCIA DAS SUBSEÇÕES

MARCA/MODELO: VXRAIL E560

FABRICANTE: DELL EMC

Este produto foi produzido de acordo com o Processo Produtivo Básico (PPB)

ITEM	QTD
VXRAIL E560, 10X2.5" ALL FLASH	2
POWEREDGE R640 MLK PLACA-MÃE	2
TRUSTED PLATFORM MODULE 2.0	2
2.5" CHASSIS COM ATÉ 10 HARD DRIVES E 3PCIE SLOTS	2
INTEL® XEON® GOLD 6230 2.1G, 20C/40T, 10.4GT/S, 27.5M CACHE, TURBO, HT (125W) DDR4-2933	2
DISSIPADOR PARA 165W	2
2933MT/S RDIMMS	2
32GB RDIMM, 2933MT/S, DUAL RANK,BCC	8
CONTROLADORA HBA330 12GBPS SAS HBA, MINICARD	2
400GB SSD SAS ESCRITA INTENSIVA 12GBPS 512N 2.5IN HOT-PLUG DRIVE, PX05SM,10 DWPD,7300 TBW	4
3.84TB SSD SATA READ INTENSIVE 6GBPS 512 2.5IN HOT-PLUG AG DRIVE, 1 DWPD, 7008 TBW	10
BOSS CONTROLLER CARD + COM 2 M.2 STICKS 240G (RAID 1),LP	2
LKEY,SW,EMC,RES,BZ,CHI, 5 ANOS	2
VXRAIL VMWARE, VSAN STANDARD, 5 ANOS	2
HCIA RECOVERPOINT PARA VMWARE PARA 1 NODE	2
IDRAC9,ENTERPRISE	2
RISER CONFIG 4, 2X16 LP	2
INTEL X550 QUATRO PORTAS 10GBE BASE-T, RNDC	2
2X 64GB MICROSDHC/SDXC CARD	2
5 STANDARD VENTILADORES PARA R640	2
CONJUNTO DE DUAS FONTES, HOT-PLUG, REDUNDANTES (1+1), 1100W	2
CABO DE FORÇA, C13, BR14136, 1.8 METROS, 250V, 10A	4
VXRAIL E560/F BEZEL	2



[www.lta-rh.com.br](http://www.lta-rh.com.br) | [comercial@lta-rh.com.br](mailto:comercial@lta-rh.com.br)

Matriz | Av. Ipiranga, 2640 | Santa Cecília | Porto Alegre | RS | Brasil | CEP 90610-000 | (51) 3382.7700/3094.1500

Filial DF | SCN, Quadra 04 | Bloco B | Sala 1251/1252 | Distrito Federal | DF | Brasil | CEP:70714-900 | (61) 3034-3004

Filial SP | Av. Paulista, 1636 | Conj.706 | Bela Vista | São Paulo | SP | Brasil | CEP 01310-200 | (11) 2391-9461

Filial RJ | Praia de Botafogo 501 | Blc I Sala 101 | Botafogo | Rio de Janeiro | RJ | CEP 22.250-040 | (21) 2586-6000

Filial PR | Rua Comendador Araújo 499 | CONJ 1007 | Centro | Curitiba | PR | CEP: 80.420-000 | (41) 99104-3240

Filial MG | Av. Do Contorno, 6594 | 705 | Belo Horizonte | MG | CEP 30110-044 | (31) 3555-3477

TRILHOS DESLIZANTES COM BRAÇO DE GERENCIAMENTO DE CABOS	2
60 MESES DE GARANTIA E SUPORTE	2
60 MESES, VSAN, STANDARD, 1 PROCESSOR-BZ	2

1.1 Fornecimento, instalação e garantia por 60 (sessenta) meses, com prestação de serviços de assistência técnica mediante manutenção corretiva da solução de hiperconvergência, com as seguintes características cada:

1.2 Solução de hiperconvergência baseada em appliances, segundo a convenção da Associação da Indústria de Redes de Armazenamento – SNIA (Storage Networking Industry Association);

1.3 A solução é composta de, no mínimo, 2 (dois) nós. Nesse caso, os servidores serão interligados diretamente (ligação direta, sem necessidade de switches) para a composição da solução de hiperconvergência.

1.4 Caso a solução não puder ser construída da forma descrita no item 1.3 e possuir alguma das características abaixo:

- Solução só funcionar com mais de dois nós e precisarem de ligação de 10Gbps para funcionar corretamente segundo melhores práticas do fabricante;
- Solução funcionar com dois nós, mas não puderem ser interligados diretamente e precisarem de ligação de 10Gbps para funcionar corretamente segundo melhores práticas do fabricante;
- Deverão ser fornecidos switches de rede e outros componentes que se façam necessários para plena operacionalização da solução, de forma que a solução seja totalmente redundante;

1.4.1 No caso de fornecimento de switches de rede, estes deverão ser gerenciáveis e possuir cada um, 1 interface UTP de 1 Gbps UTP para gerência;

1.4.2 Os switches de rede deverão ter portas suficientes para interligar todos os nós da solução;

1.5 A solução deve empregar recursos de alta disponibilidade para garantir a continuidade dos serviços, mesmo em caso de falha parcial dos equipamentos, e prevê recursos de recuperação contra desastres em caso de falha;

1.6 O sistema gerenciador da solução suporta atualizações de todos os nós da solução de forma simples e sem interrupção dos serviços para os usuários;

1.7 A solução está estruturada de forma a suportar a implementação de ambiente de virtualização em alta disponibilidade, conforme as boas práticas do fabricante da solução de virtualização;

1.8 A solução suporta e é compatível com, no mínimo, o softwares de virtualização VMWare, para consolidação de servidores, além de garantir recursos de recuperação automática em caso de falhas de hardware, evitando tempo de parada para manutenção;

1.9 No que diz respeito à disponibilidade dos dados, a solução garante que todas as informações sejam replicadas em pelo menos 1 (um) nó de maneira síncrona, assegurando a resiliência da solução e alta disponibilidade. A solução possui ferramenta de checagem interna integrada a console de gerenciamento, buscando por problemas relativos à saúde na solução proativamente;

1.10 A solução emprega portal para alocação de recursos, suportado por VMware vCenter 6 e superior;



[www.lta-rh.com.br](http://www.lta-rh.com.br) | [comercial@lta-rh.com.br](mailto:comercial@lta-rh.com.br)

**Matriz** | Av. Ipiranga, 2640 | Santa Cecília | Porto Alegre | RS | Brasil | CEP 90610-000 | (51) 3382.7700/3094.1500

**Filial DF** | SCN, Quadra 04 | Bloco B | Sala 1251/1252 | Distrito Federal | DF | Brasil | CEP:70714-900 | (61) 3034-3004


**Filial SP** | Av. Paulista, 1636 | Conj.706 | Bela Vista | São Paulo | SP | Brasil | CEP 01310-200 | (11) 2391-9461

**Filial RJ** | Praia de Botafogo 501 | Blc I Sala 101 | Botafogo | Rio de Janeiro | RJ | CEP 22.250-040 | (21) 2586-6000

**Filial PR** | Rua Comendador Araújo 499 | CONJ 1007 | Centro | Curitiba | PR | CEP: 80.420-000 | (41) 99104-3240

**Filial MG** | Av. Do Contorno, 6594 | 705 | Belo Horizonte | MG | CEP 30110-044 | (31) 3555-3477

- 1.11 Toda a solução tem suporte centralizado em uma única central de atendimento do fornecedor da solução de hardware;
- 1.12 Todos os manuais técnicos referentes aos componentes da solução serão fornecidos ou disponibilizados eletronicamente;
- 1.13 Cada módulo possui, no mínimo, 4 (quatro) interfaces de rede padrão GbE/10GbE (1000BASE-T /10GBASE-T) para interligação à rede de datacenter das localidades por meio de cabos padrão UTP;
- 1.14 A solução permite instalação, em rack padrão 19 (dezenove) polegadas:
- 1.14.1 Cada nó de hiperconvergência ocupa, no máximo 2 U (duas unidades de rack);
- 1.14.2 Os nós possuem botão liga/desliga com proteção para prevenir o desligamento acidental;
- 1.14.3 Possui display ou led frontal para exibição de alertas de funcionamento dos componentes internos, tais como falhas de processador, memória RAM, fontes de alimentação, disco rígido e ventilador;
- 1.14.4 Será entregue junto com cada módulo o correspondente kit de fixação no rack, do tipo retrátil, permitindo o deslizamento do servidor a fim de facilitar sua manutenção;
- 1.15 Possui sistema de ventilação redundante e hotpluggable para que a CPU suporte à configuração máxima e dentro dos limites de temperatura adequados ao perfeito funcionamento do equipamento, permitindo a substituição mesmo com o equipamento em funcionamento;
- 1.16 A solução possui, no mínimo, 2 (duas) fontes, suportando o funcionamento integral do equipamento mesmo em caso de falha de uma das fontes;
- 1.17 As fontes são redundantes e hot-pluggable, permitindo a substituição de qualquer uma das fontes em caso de falha sem parada ou comprometimento do funcionamento do equipamento:
- 1.17.1 Cada uma das fontes tem potência mínima suficiente para atender adequadamente a configuração máxima do equipamento entregue;
- 1.17.2 As fontes possuem tensão de entrada de 100VAC a 240VAC a 60Hz, com ajuste automático de tensão;
- 1.17.3 Acompanha cabo de alimentação para cada fonte de alimentação fornecida;
- 1.18 A solução emprega as seguintes configurações:
- 1.18.1 Capacidade total líquida de processamento da solução é de, no mínimo, 10 núcleos (retirando-se a carga utilizada para operação e gerenciamento da solução) considerando inclusive cenário de operação parcial da solução (falha de algum nó);
- 1.18.2 Possuir suporte a hyperthread, frequência de clock interno de, no mínimo, 2.10 GHz;
- 1.18.3 A solução entregue aumenta progressivamente a quantidade de núcleos total fornecida de maneira proporcional à quantidade de recursos demandados pelos componentes relativos à infraestrutura interna de controle e gerenciamento da solução operando nos nós (exemplo: máquinas virtuais para administração do armazenamento SDS - Software-Defined Storage);
- 1.18.4 Cada nó da solução possui controladora de memória com suporte a DDR4 de, no mínimo, 2400MHz e 6 canais de memória;
- 1.19 A solução opera com pentes de memória RAM tipo DDR4 RDIMM (Registered DIMM) ou LRDIMM (Load Reduced DIMM), com tecnologia de correção ECC (Error Correcting Code) e velocidade de, no mínimo, 2400MHz, e possui as seguintes características:
- 1.19.1 A solução entrega uma capacidade líquida mínima de memória RAM (retirando-se o montante utilizado para operação da solução) de 96 GB considerando inclusive cenário de operação parcial da solução (falha de algum nó);



[www.lta-rh.com.br](http://www.lta-rh.com.br) | [comercial@lta-rh.com.br](mailto:comercial@lta-rh.com.br)

**Matriz** | Av. Ipiranga, 2640 | Santa Cecília | Porto Alegre | RS | Brasil | CEP 90610-000 | (51) 3382.7700/3094.1500

**Filial DF** | SCN, Quadra 04 | Bloco B | Sala 1251/1252 | Distrito Federal | DF | Brasil | CEP:70714-900 | (61) 3034-3004

**Filial SP** | Av. Paulista, 1636 | Conj.706 | Bela Vista | São Paulo | SP | Brasil | CEP 01310-200 | (11) 2391-9461

**Filial RJ** | Praia de Botafogo 501 | Blc I Sala 101 | Botafogo | Rio de Janeiro | RJ | CEP 22.250-040 | (21) 2586-6000

**Filial PR** | Rua Comendador Araújo 499 | CONJ 1007 | Centro | Curitiba | PR | CEP: 80.420-000 | (41) 99104-3240

**Filial MG** | Av. Do Contorno, 6594 | 705 | Belo Horizonte | MG | CEP 30110-044 | (31) 3555-3477

- 1.20 Os recursos de armazenamento são compartilhados entre todos os servidores e appliances da solução por meio de armazenamento definido por software - SDS;
- 1.21 A solução entrega uma capacidade líquida de armazenamento de, no mínimo, 12 TB, considerando inclusive cenário de operação parcial da solução (falha de algum nó);
- 1.21.1 A tecnologia de discos de armazenamento da solução poderá ser composta de qualquer tecnologia de discos, de forma híbrida ou não, desde que não seja totalmente composta somente por discos NL-SAS;
- 1.21.2 No caso da composição dos discos dos nós da solução ser composta por discos NL-SAS, deve-se respeitar a proporção de, no mínimo, 30% de armazenamento líquido em discos SSD;
- 1.22 É permitida a troca de nós avariados, sem interrupção das operações e I/O das aplicações que acessarão os dados;
- 1.23 A falha isolada de um nó da solução não poderá impactar a disponibilidade da infraestrutura de armazenamento para as máquinas virtuais;
- 1.24 A solução permite a operação em grupo de, pelo menos, 2 appliances do mesmo tipo;
- 1.25 A solução possui console de administração web que permite integração com Microsoft Active Directory para autenticação;
- 1.26 Provê monitoração automática e periódica da solução, com o envio de notificações preventivamente em caso de falhas, notificando o suporte do fornecedor a tomar medidas preventivas e acordadas com o CONTRATANTE a fim de evitar tempo de inatividade e impactos na produção;
- 1.27 Oferece portal de acesso do próprio fornecedor para download de atualizações e de softwares agregados à solução a fim de atender rapidamente demandas dos negócios;
- 1.28 A ferramenta de gerenciamento deverá fornecer um dashboard reportando a utilização dos recursos da solução como CPU, memória e armazenamento;
- 1.29 Comprovar que o equipamento não é fabricado utilizando as seguintes substâncias nocivas ao meio ambiente nas quantidades permitidas pela diretiva RoHS: chumbo (Pb), cádmio (Cd), mercúrio (Hg), cromo hexavalente (Hex-CR, bifenilos polibromados (PBBs) e éteres defenil-polibromados (PBDEs);

## 2 ITEM 2 - INSTALAÇÃO DA SOLUÇÃO DE HIPERCONVERGÊNCIA

- 2.1 A instalação da solução de hiperconvergência deverá ocorrer nos locais conforme distribuição no item 05 do Anexo I deste Edital I;
- 2.2 Todo ferramental necessário para execução dos serviços de instalação, configuração inicial, incluindo softwares, máquinas virtuais, equipamentos ou ferramentas, bem como eventuais materiais necessários para ligações temporárias, são de inteira responsabilidade da CONTRATADA;
- 2.3 A CONTRATANTE disponibilizará o espaço no CPD, refrigeração suficiente para comportar a solução a ser adquirida, assim como, a infraestrutura elétrica até a posição onde será instalado a solução, com capacidades (corrente e tensão) suficientes de suportar todos os equipamentos.
- 2.4 Entende-se por instalação, para efeito deste projeto:
- 2.4.1 Instalação física e lógica do equipamento, com fornecimento de cabos, UTP Cat. 6 para a totalidade das portas da solução e de gerenciamento, interligando os mesmos às redes LAN;
- 2.4.2 A metragem dos cabos fornecidos devem ser de, no mínimo, 5 metros;
- 2.4.3 As cores dos cabos fornecidos devem ser:
- 2.4.3.1 4 (quatro) cabos azuis;
- 2.4.3.2 1 (um) cabo cinza (a ser usado para gerenciamento do servidor);
- 2.5 A solução será instalada na última versão de firmware disponível pelo fabricante;



[www.lta-rh.com.br](http://www.lta-rh.com.br) | [comercial@lta-rh.com.br](mailto:comercial@lta-rh.com.br)

**Matriz** | Av. Ipiranga, 2640 | Santa Cecília | Porto Alegre | RS | Brasil | CEP 90610-000 | (51) 3382.7700/3094.1500

**Filial DF** | SCN, Quadra 04 | Bloco B | Sala 1251/1252 | Distrito Federal | DF | Brasil | CEP:70714-900 | (61) 3034-3004

**Filial SP** | Av. Paulista, 1636 | Conj.706 | Bela Vista | São Paulo | SP | Brasil | CEP 01310-200 | (11) 2391-9461

**Filial RJ** | Praia de Botafogo 501 | Blc I Sala 101 | Botafogo | Rio de Janeiro | RJ | CEP 22.250-040 | (21) 2586-6000

**Filial PR** | Rua Comendador Araújo 499 | CONJ 1007 | Centro | Curitiba | PR | CEP: 80.420-000 | (41) 99104-3240

**Filial MG** | Av. Do Contorno, 6594 | 705 | Belo Horizonte | MG | CEP 30110-044 | (31) 3555-3477

2.6 A instalação do equipamento será executada pela contratada, prestados por profissionais devidamente qualificados e/ou diretamente pelo fabricante dos equipamentos;

2.7 Entende-se por configuração inicial, para efeito deste projeto:

2.7.1 Elaboração em conjunto com a equipe técnica do CONTRATANTE, de projeto de configuração, segundo as melhores práticas do fabricante e considerando as demandas e características dos serviços do CONTRATANTE;

2.7.2 Realização da configuração inicial da solução ofertada, segundo projeto, e conforme padrão de endereçamento IP a ser fornecido pelo CONTRATANTE;

2.7.3 Migração de uma máquina virtual existente no ambiente do CONTRATANTE para a solução ofertada e realização de teste de falha ao simular falha de um nó sem acarretar indisponibilidade do serviço provido por essa máquina virtual.

2.8 Fornecimento e configuração de software de SDS padrão do fabricante;



[www.lta-rh.com.br](http://www.lta-rh.com.br) | [comercial@lta-rh.com.br](mailto:comercial@lta-rh.com.br)

**Matriz** | Av. Ipiranga, 2640 | Santa Cecília | Porto Alegre | RS | Brasil | CEP 90610-000 | (51) 3382.7700/3094.1500

**Filial DF** | SCN, Quadra 04 | Bloco B | Sala 1251/1252 | Distrito Federal | DF | Brasil | CEP:70714-900 | (61) 3034-3004

**Filial SP** | Av. Paulista, 1636 | Conj.706 | Bela Vista | São Paulo | SP | Brasil | CEP 01310-200 | (11) 2391-9461

**Filial RJ** | Praia de Botafogo 501 | Blc I Sala 101 | Botafogo | Rio de Janeiro | RJ | CEP 22.250-040 | (21) 2586-6000

**Filial PR** | Rua Comendador Araújo 499 | CONJ 1007 | Centro | Curitiba | PR | CEP: 80.420-000 | (41) 99104-3240

**Filial MG** | Av. Do Contorno, 6594 | 705 | Belo Horizonte | MG | CEP 30110-044 | (31) 3555-3477



**ANEXO II - PREGÃO ELETRÔNICO SRP Nº 40/2019  
PLANILHA PARA FORMULAÇÃO DE PREÇOS**

**Ao**  
**TRIBUNAL REGIONAL FEDERAL DA PRIMEIRA REGIÃO – TRF 1ª REGIÃO**  
**Ref.: PREGÃO ELETRÔNICO RP Nº 40/2019**  
**Proposta N.º 281/19**

GRUPO	ITEM	DESCRIÇÃO	UN.	QUAN T.	VALOR UNITÁRIO (R\$)	VALOR TOTAL (R\$)
1	01	Solução de hiperconvergência	Unidade	82	R\$ 155.000,00	R\$ 12.710.000,00
	02	Instalação – Solução de hiperconvergência	Serviço	82	R\$ 4.000,00	R\$ 328.000,00
<b>VALOR TOTAL DO GRUPO 01</b>					<b>R\$ 13.038.000,00</b>	

**OBSERVAÇÕES:**

- a) prazo de entrega dos equipamentos, de 60 (sessenta) dias corridos, contados do recebimento da Ordem de Fornecimento e prazo de execução dos serviços de instalação, configuração e teste, que não poderá ser superior a 10 (dez) dias úteis, contados do recebimento da Ordem de execução dos Serviços;
- b) prazo de assistência técnica da garantia dos equipamentos, de 60 (sessenta) meses, contados a partir da data de assinatura do Termo de Recebimento Definitivo;
- c) validade da proposta, de 60 (sessenta) dias, contados do dia útil imediatamente posterior ao indicado no item 2 do Edital;

**Declaramos que não estamos sob pena de interdição de direitos previstos na Lei nº 9.605, de 12.02.98 (Lei de Crimes Ambientais);**

**OUTRAS OBSERVAÇÕES:**

1 - Além das condições constantes do subitem 5.2 deste Edital, deverão constar da proposta os seguintes dados do REPRESENTANTE LEGAL que assinará o Contrato:

**A) NOME COMPLETO:** Alexander Costa Barcelos

**B) CARTEIRA DE IDENTIDADE:** 2035263058

**C) CPF:** 594.509.830-20

**D) E-MAIL:** [comercial@lta-rh.com.br](mailto:comercial@lta-rh.com.br)

**E) TELEFONE:** (51) 3382.7700

**F) CELULAR:** (51) 99597-0360

**G) DOMICÍLIO:** Rua Farnese nº 110/401, Bela Vista, Porto Alegre – RS. – cep 90450-180

**DATA E HORA DA SESSÃO PÚBLICA:** 10/10/2019 ÀS 14:00.

Brasília, 1º de novembro de 2019.

\_\_\_\_\_  
ALEXANDER C. BARCELOS  
DIRETOR COMERCIAL  
CPF: 594.509.830-20 | RG: 2035263058



**www.lta-rh.com.br | comercial@lta-rh.com.br**

**Matriz** | Av. Ipiranga, 2640 | Santa Cecília | Porto Alegre | RS | Brasil | CEP 90610-000 | (51) 3382.7700/3094.1500

**Filial DF** | SCN, Quadra 04 | Bloco B | Sala 1251/1252 | Distrito Federal | DF | Brasil | CEP:70714-900 | (61) 3034-3004

**Filial SP** | Av. Paulista, 1636 | Conj.706 | Bela Vista | São Paulo | SP | Brasil | CEP 01310-200 | (11) 2391-9461

**Filial RJ** | Praia de Botafogo 501 | Blc I Sala 101 | Botafogo | Rio de Janeiro | RJ | CEP 22.250-040 | (21) 2586-6000

**Filial PR** | Rua Comendador Araújo 499 | CONJ 1007 | Centro | Curitiba | PR | CEP: 80.420-000 | (41) 99104-3240

**Filial MG** | Av. Do Contorno, 6594 | 705 | Belo Horizonte | MG | CEP 30110-044 | (31) 3555-3477

FORMULÁRIO DE AVALIAÇÃO TÉCNICA

Ao  
 TRIBUNAL REGIONAL FEDERAL DA PRIMEIRA REGIÃO – TRF 1ª REGIÃO  
 Ref.: PREGÃO ELETRÔNICO RP Nº 40/2019  
 Proposta N.º 281/19

1. O formulário a partir do modelo constante do presente anexo é de preenchimento obrigatório, e deverá fazer parte integrante da proposta técnica de cada licitante;
2. As propostas que não atenderem à totalidade das características obrigatórias serão desclassificadas;
3. O formulário deverá ser preenchido sob a seguinte orientação:
  - a) Coluna Item no anexo I-A: indicar a página da proposta da licitante que contém documento que comprove a característica solicitada;
  - b) Coluna Item da Proposta: indicar o item, subitem, parágrafo, alínea, etc, do documento que comprove a característica solicitada.
  - c) Coluna Página do Catálogo/folder/manual: indicar a página da proposta da licitante que contém documento que comprove a característica solicitada;

ITEM DO EDITAL E DA ESPECIFICAÇÃO TÉCNICA	AVALIAÇÃO TÉCNICA DA PROPOSTA		
	ITEM NO ANEXO I-A Página da proposta da licitante que contém documento que comprove a característica solicitada	ITEM DA PROPOSTA Item, subitem, parágrafo, alínea, etc, do documento que comprove a característica solicitada.	PÁGINA DO CATÁLOGO/FOLDER/M ANUAL Página da proposta da licitante que contém documento que comprove a característica solicitada
<b>ITEM 1 - SOLUÇÃO DE HIPERCONVERGÊNCIA DAS SUBSEÇÕES</b>			
1.1 Fornecimento, instalação e garantia por 60 (sessenta) meses, com prestação de serviços de assistência técnica mediante manutenção corretiva da solução de hiperconvergência, com as seguintes características cada:	De acordo, atendemos ao solicitado, conforme consta nesta proposta		
1.2 Solução de hiperconvergência baseada em appliances, segundo a convenção da Associação da Indústria de Redes de Armazenamento – SNIA	Pag. 9 (VXRAIL E560 APPLIANCE)	Item: APPLIANCE DELL EMC VXRAIL  <b>Declaração técnica</b>	Página 01 do Catálogo "h16763-vxrail-14g- spec-sheet"



[www.lta-rh.com.br](http://www.lta-rh.com.br) | [comercial@lta-rh.com.br](mailto:comercial@lta-rh.com.br)

**Matriz** | Av. Ipiranga, 2640 | Santa Cecília | Porto Alegre | RS | Brasil | CEP 90610-000 | (51) 3382.7700/3094.1500  
**Filial DF** | SCN, Quadra 04 | Bloco B | Sala 1251/1252 | Distrito Federal | DF | Brasil | CEP:70714-900 | (61) 3034-3004  
**Filial SP** | Av. Paulista, 1636 | Conj.706 | Bela Vista | São Paulo | SP | Brasil | CEP 01310-200 | (11) 2391-9461  
**Filial RJ** | Praia de Botafogo 501 | Blc I Sala 101 | Botafogo | Rio de Janeiro | RJ | CEP 22.250-040 | (21) 2586-6000  
**Filial PR** | Rua Comendador Araújo 499 | CONJ 1007 | Centro | Curitiba | PR | CEP: 80.420-000 | (41) 99104-3240  
**Filial MG** | Av. Do Contorno, 6594 | 705 | Belo Horizonte | MG | CEP 30110-044 | (31) 3555-3477

<p>(Storage Networking Industry Association);</p>		<p>Item: Solução de hiperconvergência baseada em appliances, segundo a convenção da Associação da Indústria de Redes de Armazenamento – SNIA (Storage Networking Industry Association) - SNIA (Storage Networking Industry Association);</p>	<p>Página 01 da <b>Declaração técnica</b></p>
<p>1.3 A solução deverá ser composta de, no mínimo, 2 (dois) nós. Nesse caso, os servidores deverão ser interligados diretamente (ligação direta, sem necessidade de switches) para a composição da solução de hiperconvergência.</p>	<p>Pag. 9 (VXRAIL E560 qtd: 2)</p>	<p>Item: Requirements, Recommendations, and Restrictions Subitem: PHYSICAL NETWORK</p>	<p>Página 6 e 7 do catálogo "h17566-vsant-2node-cluster-on-vxrail-planning-guide"</p>
<p>1.4 Caso a solução não puder ser construída da forma descrita no item 1.3 e possuir alguma das características abaixo:</p> <ul style="list-style-type: none"> <li>· Solução só funcionar com mais de dois nós e precisarem de ligação de 10Gbps para funcionar corretamente segundo melhores práticas do fabricante;</li> <li>· Solução funcionar com dois nós, mas não puderem ser interligados diretamente e precisarem de ligação de 10Gbps para funcionar corretamente segundo melhores práticas do</li> </ul>	<p><b>Não se aplica à nossa oferta</b></p>		



[www.lta-rh.com.br](http://www.lta-rh.com.br) | [comercial@lta-rh.com.br](mailto:comercial@lta-rh.com.br)

**Matriz** | Av. Ipiranga, 2640 | Santa Cecília | Porto Alegre | RS | Brasil | CEP 90610-000 | (51) 3382.7700/3094.1500

**Filial DF** | SCN, Quadra 04 | Bloco B | Sala 1251/1252 | Distrito Federal | DF | Brasil | CEP:70714-900 | (61) 3034-3004

**Filial SP** | Av. Paulista, 1636 | Conj.706 | Bela Vista | São Paulo | SP | Brasil | CEP 01310-200 | (11) 2391-9461

**Filial RJ** | Praia de Botafogo 501 | Blc I Sala 101 | Botafogo | Rio de Janeiro | RJ | CEP 22.250-040 | (21) 2586-6000

**Filial PR** | Rua Comendador Araújo 499 | CONJ 1007 | Centro | Curitiba | PR | CEP: 80.420-000 | (41) 99104-3240

**Filial MG** | Av. Do Contorno, 6594 | 705 | Belo Horizonte | MG | CEP 30110-044 | (31) 3555-3477

<p>fabricante; · Deverão ser fornecidos switches de rede e outros componentes que se façam necessários para plena operacionalização da solução, de forma que a solução seja totalmente redundante;</p>			
<p>1.4.1 No caso de fornecimento de switches de rede, estes deverão ser gerenciáveis e possuir cada um, 1 interface UTP de 1 Gbps UTP para gerência;</p>	<p><b>Não se aplica à nossa oferta</b></p>		
<p>1.4.2 Os switches de rede deverão ter portas suficientes para interligar todos os nós da solução;</p>	<p><b>Não se aplica à nossa oferta</b></p>		
<p>1.5 A solução deve empregar recursos de alta disponibilidade para garantir a continuidade dos serviços, mesmo em caso de falha parcial dos equipamentos, e deve prever recursos de recuperação contra desastres em caso de falha;</p>	<p>Pag. 9 (VXRAIL E560 qtd: 2)</p>	<p>Item: Failures to Tolerate Subitem: This FTT option generally defines the number of host and device failures that a virtual machineobject can tolerate item: 2-node configuration subitem: The configuration only</p>	<p>Pagina 63 e 76 do catálogo "h15104-vxrail-appliance-techbook"</p>



[www.lta-rh.com.br](http://www.lta-rh.com.br) | [comercial@lta-rh.com.br](mailto:comercial@lta-rh.com.br)

**Matriz** | Av. Ipiranga, 2640 | Santa Cecília | Porto Alegre | RS | Brasil | CEP 90610-000 | (51) 3382.7700/3094.1500

**Filial DF** | SCN, Quadra 04 | Bloco B | Sala 1251/1252 | Distrito Federal | DF | Brasil | CEP:70714-900 | (61) 3034-3004

**Filial SP** | Av. Paulista, 1636 | Conj.706 | Bela Vista | São Paulo | SP | Brasil | CEP 01310-200 | (11) 2391-9461

**Filial RJ** | Praia de Botafogo 501 | Blc I Sala 101 | Botafogo | Rio de Janeiro | RJ | CEP 22.250-040 | (21) 2586-6000

**Filial PR** | Rua Comendador Araújo 499 | CONJ 1007 | Centro | Curitiba | PR | CEP: 80.420-000 | (41) 99104-3240

**Filial MG** | Av. Do Contorno, 6594 | 705 | Belo Horizonte | MG | CEP 30110-044 | (31) 3555-3477

		supports mirroring (FTT=1)	
1.6 O sistema gerenciador da solução deverá suportar atualizações de todos os nós da solução de forma simples e sem interrupção dos serviços para os usuários;	Pag. 9 (VXRAIL SOFTWARE)	Item: VxRail software architecture Subitem: VxRail Manager	Página 34 do catálogo "h15104-vxrail-appliance-techbook"
1.7 A solução deverá estar estruturada de forma a suportar a implementação de ambiente de virtualização em alta disponibilidade, conforme as boas práticas do fabricante da solução de virtualização;	Pag. 9 (VMWARE VSAN)	<b>Declaração técnica</b> Item: A solução está estruturada de forma a suportar a implementação de ambiente de virtualização em alta disponibilidade, conforme as boas práticas do fabricante da solução de virtualização;	Página 01 da <b>Declaração técnica</b>
1.8 A solução deverá suportar e ser compatível com, no mínimo, o softwares de virtualização VMWare, para consolidação de servidores, além de garantir recursos de recuperação automática em caso de falhas de hardware, evitando tempo de parada para manutenção;	Pag. 9 (VMWARE VSAN)	Item: VxRail software architecture Subitem: vSphere and vSAN ordering information	Página 38 e 39 do catálogo "h15104-vxrail-appliance-techbook"



[www.lta-rh.com.br](http://www.lta-rh.com.br) | [comercial@lta-rh.com.br](mailto:comercial@lta-rh.com.br)

**Matriz** | Av. Ipiranga, 2640 | Santa Cecília | Porto Alegre | RS | Brasil | CEP 90610-000 | (51) 3382.7700/3094.1500

**Filial DF** | SCN, Quadra 04 | Bloco B | Sala 1251/1252 | Distrito Federal | DF | Brasil | CEP:70714-900 | (61) 3034-3004

**Filial SP** | Av. Paulista, 1636 | Conj.706 | Bela Vista | São Paulo | SP | Brasil | CEP 01310-200 | (11) 2391-9461

**Filial RJ** | Praia de Botafogo 501 | Blc I Sala 101 | Botafogo | Rio de Janeiro | RJ | CEP 22.250-040 | (21) 2586-6000

**Filial PR** | Rua Comendador Araújo 499 | CONJ 1007 | Centro | Curitiba | PR | CEP: 80.420-000 | (41) 99104-3240

**Filial MG** | Av. Do Contorno, 6594 | 705 | Belo Horizonte | MG | CEP 30110-044 | (31) 3555-3477



<p>1.9 No que diz respeito à disponibilidade dos dados, a solução deverá garantir que todas as informações sejam replicadas em pelo menos 1 (um) nó de maneira síncrona, assegurando a resiliência da solução e alta disponibilidade. A solução deverá possuir ferramenta de checagem interna integrada a console de gerenciamento, buscando por problemas relativos à saúde na solução proativamente;</p>	<p>Pag. 9 (VMWARE VSAN)</p>	<p>Item: Failures to Tolerate Subitem: This FTT option generally defines the number of host and device failures that a virtual machineobject can tolerate item: 2-node configuration subitem: The configuration only supports mirroring (FTT=1)</p>	<p>Página 63 e 76 do catálogo "h15104-vxrail-appliance-techbook"</p>
<p>1.10 A solução deverá empregar portal para alocação de recursos, suportado por VMware vCenter 6 e superior;</p>	<p>Pag. 9 (VXRAIL SOFTWARE)</p>	<p>Item: Dell EMC VxRail Systems subitem: Managed through the ubiquitous VMware vCenter Server interface, VxRail provides a familiar vSphere experience</p>	<p>Página 15 do catálogo "h15104-vxrail-appliance-techbook"</p>
<p>1.11 Toda a solução deverá ter suporte centralizado em uma única central de atendimento do fornecedor da solução de hardware;</p>	<p>Pag. 10 (Suporte de software e hardware)</p>	<p><b>Declaração técnica</b> Item: Toda a solução possui suporte centralizado em uma única central de atendimento do fornecedor da solução de hardware;</p>	<p>Página 01 da <b>Declaração técnica</b></p>
<p>1.12 Todos os manuais técnicos referentes aos componentes da solução deverão ser fornecidos ou disponibilizados eletronicamente;</p>	<p>De acordo, atendemos ao solicitado, conforme consta nesta proposta</p>		
<p>1.13 Cada módulo deverá possuir, no mínimo, 4 (quatro) interfaces de rede padrão GbE/10GbE (1000BASE-T /10GBASE-T)</p>	<p>Pag. 10 (INTEL X550 QUATRO Portas 10GBE BASE-T, RNDG)</p>	<p>Item: Ports and connectors specifications Subitem: NIC ports</p>	<p>Página 20 do catálogo "docu87343_VxRail-Appliances-on-14th-Generation-PowerEdge-Servers-E-Series-</p>



[www.lta-rh.com.br](http://www.lta-rh.com.br) | [comercial@lta-rh.com.br](mailto:comercial@lta-rh.com.br)

**Matriz** | Av. Ipiranga, 2640 | Santa Cecília | Porto Alegre | RS | Brasil | CEP 90610-000 | (51) 3382.7700/3094.1500  
**Filial DF** | SCN, Quadra 04 | Bloco B | Sala 1251/1252 | Distrito Federal | DF | Brasil | CEP:70714-900 | (61) 3034-3004  
**Filial SP** | Av. Paulista, 1636 | Conj.706 | Bela Vista | São Paulo | SP | Brasil | CEP 01310-200 | (11) 2391-9461  
**Filial RJ** | Praia de Botafogo 501 | Blc I Sala 101 | Botafogo | Rio de Janeiro | RJ | CEP 22.250-040 | (21) 2586-6000  
**Filial PR** | Rua Comendador Araújo 499 | CONJ 1007 | Centro | Curitiba | PR | CEP: 80.420-000 | (41) 99104-3240  
**Filial MG** | Av. Do Contorno, 6594 | 705 | Belo Horizonte | MG | CEP 30110-044 | (31) 3555-3477

para interligação à rede de datacenter das localidades por meio de cabos padrão UTP;			Owner's-Manual"
1.14 A solução deverá permitir instalação, em rack padrão 19 (dezenove) polegadas:	Pag. 10 (TRILHOS DESLIZANTES COM BRAÇO DE GERENCIAMENTO DE CABOS)	Item: CHASSIS Subitem: 1U1N	Página 2 do catálogo "h16763-vxrail-14g-spec-sheet.pdf"
1.14.1 Cada nó de hiperconvergência deverá ocupar, no máximo 2 U (duas unidades de rack);	Pag. 9 (VXRAIL E560)	Item: VxRail hardware architecture Subitem: E Series	Página 19 do catálogo "h15104-vxrail-appliance-techbook"
1.14.2 Os nós deverão possuir botão liga/desliga com proteção para prevenir o desligamento acidental;	Pag. 9 (VXRAIL E560)	Item: Overview Subitem: Right control panel view  <b>Declaração Técnica</b> Item: Possui botão liga/desliga com proteção para prevenir o desligamento acidental	Página 8 do catálogo "docu87343_VxRail-Appliances-on-14th-Generation-PowerEdge-Servers-E-Series-Owner's-Manual"  Página 01 da <b>Declaração Técnica</b>
1.14.3 Deverá possuir display ou led frontal para exibição de alertas de funcionamento dos componentes internos, tais como falhas de processador, memória RAM, fontes de alimentação, disco rígido e ventilador;	Pag. 9 (VXRAIL E560)	Item: Front view Subitem: Status LED indicators e Appliance health and appliance ID indicator codes	Página 6 e 7 do catálogo "docu87343_VxRail-Appliances-on-14th-Generation-PowerEdge-Servers-E-Series-Owner's-Manual"
1.14.4 Deverá ser entregue junto com cada módulo o correspondente kit de fixação no rack, do tipo retrátil, permitindo o deslizamento do servidor a fim de facilitar sua manutenção;	Pag. 9 (VXRAIL E560)	<b>Declaração técnica</b> Item: Será entregue junto com cada módulo kit de fixação no rack, do tipo retrátil, permitindo o deslizamento do servidor a fim de facilitar sua manutenção;	Página 01 da <b>Declaração técnica</b>



[www.lta-rh.com.br](http://www.lta-rh.com.br) | [comercial@lta-rh.com.br](mailto:comercial@lta-rh.com.br)

**Matriz** | Av. Ipiranga, 2640 | Santa Cecília | Porto Alegre | RS | Brasil | CEP 90610-000 | (51) 3382.7700/3094.1500

**Filial DF** | SCN, Quadra 04 | Bloco B | Sala 1251/1252 | Distrito Federal | DF | Brasil | CEP:70714-900 | (61) 3034-3004

**Filial SP** | Av. Paulista, 1636 | Conj.706 | Bela Vista | São Paulo | SP | Brasil | CEP 01310-200 | (11) 2391-9461

**Filial RJ** | Praia de Botafogo 501 | Blc I Sala 101 | Botafogo | Rio de Janeiro | RJ | CEP 22.250-040 | (21) 2586-6000

**Filial PR** | Rua Comendador Araújo 499 | CONJ 1007 | Centro | Curitiba | PR | CEP: 80.420-000 | (41) 99104-3240

**Filial MG** | Av. Do Contorno, 6594 | 705 | Belo Horizonte | MG | CEP 30110-044 | (31) 3555-3477

<p>1.15 Deverá possuir sistema de ventilação redundante e hotpluggable para que a CPU suporte à configuração máxima e dentro dos limites de temperatura adequados ao perfeito funcionamento do equipamento, permitindo a substituição mesmo com o equipamento em funcionamento;</p>	<p>Pag. 9 (VXRAIL E560)</p>	<p>Item: O padrão em infraestrutura hiperconvergente Subitem: Série E</p>	<p>Página 3 do catálogo "h16763-vxrail-14g-spec-sheet" Página 01 da <b>Declaração técnica</b></p>
<p>1.16 A solução deverá possuir, no mínimo, 2 (duas) fontes, suportando o funcionamento integral do equipamento mesmo em caso de falha de uma das fontes;</p>	<p>Pag. 10 (CONJUNTO DE DUAS FONTES, Hot-plug, Redundantes (1+1), 1100W)</p>	<p>Item: Overview Subitem: Back view of the appliance</p>	<p>Página 5 do catálogo "docu87343_VxRail-Appliances-on-14th-Generation-PowerEdge-Servers-E-Series-Owner's-Manual"</p>
<p>1.17 As fontes deverão ser redundantes e hot-pluggable, permitindo a substituição de qualquer uma das fontes em caso de falha sem parada ou comprometimento do funcionamento do equipamento:</p>	<p>Pag. 10 (CONJUNTO DE DUAS FONTES, Hot-plug, Redundantes (1+1), 1100W)</p>	<p>Item: Overview Subitem: Back view of the appliance</p>	<p>Página 5 do catálogo "docu87343_VxRail-Appliances-on-14th-Generation-PowerEdge-Servers-E-Series-Owner's-Manual"  Página 9 do catálogo "docu87343_VxRail-Appliances-on-14th-Generation-PowerEdge-Servers-E-Series-Owner's-Manual"</p>
<p>1.17.1 Cada uma das fontes deverá ter potência mínima suficiente para atender adequadamente a configuração máxima do equipamento entregue;</p>	<p>Pag. 10 (CONJUNTO DE DUAS FONTES, Hot-plug, Redundantes (1+1), 1100W)</p>	<p><b>Declaração técnica</b> Item: Cada fontes têm potência suficiente para atender adequadamente a configuração máxima do equipamento entregue;</p>	<p>Página 01 da <b>Declaração técnica</b></p>



[www.lta-rh.com.br](http://www.lta-rh.com.br) | [comercial@lta-rh.com.br](mailto:comercial@lta-rh.com.br)

**Matriz** | Av. Ipiranga, 2640 | Santa Cecília | Porto Alegre | RS | Brasil | CEP 90610-000 | (51) 3382.7700/3094.1500

**Filial DF** | SCN, Quadra 04 | Bloco B | Sala 1251/1252 | Distrito Federal | DF | Brasil | CEP:70714-900 | (61) 3034-3004

**Filial SP** | Av. Paulista, 1636 | Conj.706 | Bela Vista | São Paulo | SP | Brasil | CEP 01310-200 | (11) 2391-9461

**Filial RJ** | Praia de Botafogo 501 | Blc I Sala 101 | Botafogo | Rio de Janeiro | RJ | CEP 22.250-040 | (21) 2586-6000

**Filial PR** | Rua Comendador Araújo 499 | CONJ 1007 | Centro | Curitiba | PR | CEP: 80.420-000 | (41) 99104-3240

**Filial MG** | Av. Do Contorno, 6594 | 705 | Belo Horizonte | MG | CEP 30110-044 | (31) 3555-3477

<p>1.17.2 As fontes deverão possuir tensão de entrada de 100VAC a 240VAC a 60Hz, com ajuste automático de tensão;</p>	<p>Pag. 10 (CONJUNTO DE DUAS FONTES, Hot-plug, Redundantes (1+1), 1100W)</p>	<p>Item: Technical specifications Subitem: PSU specifications</p>	<p>Página 18 do catálogo "docu87343_VxRail-Appliances-on-14th-Generation-PowerEdge-Servers-E-Series-Owner's-Manual"</p>
<p>1.17.3 Deverá acompanhar cabo de alimentação para cada fonte de alimentação fornecida;</p>	<p>Pag. 10 CABO DE FORÇA, C13, BR14136, 1.8 metros, 250V, 10A</p>	<p><b>Declaração técnica</b> Item: Cada fontes têm potência suficiente para atender adequadamente a configuração máxima do equipamento entregue, <b>acompanhado de cabo de alimentação;</b></p>	<p>Página 01 da <b>Declaração técnica</b></p>
<p>1.18 A solução deverá empregar as seguintes configurações: 1.18.1 Capacidade total líquida de processamento da solução deverá ser de, no mínimo, 10 núcleos (retirando-se a carga utilizada para operação e gerenciamento da solução) considerando inclusive cenário de operação parcial da solução (falha de algum nó);</p>	<p>Pag. 9 <b>INTEL® XEON® GOLD 6230 2.1G, 20C/40T, 10.4GT/S, 27.5M CACHE, TURBO, HT (125W) DDR4-2933</b></p>	<p>Processador Intel® Xeon® Gold 6230 (cache de 27,5 M, 2,10 GHz) Product Specifications.pdf</p>	<p>Página 1 do catálogo O processador ofertado possui 20 cores à 2,1GHz portanto atende ao requisitado</p>
<p>1.18.2 Possuir suporte a hyperthread, frequência de clock interno de, no mínimo, 2.10 GHz;</p>	<p>Pag. 9 <b>INTEL® XEON® GOLD 6230 2.1G, 20C/40T, 10.4GT/S, 27.5M CACHE, TURBO, HT (125W) DDR4-2933</b></p>	<p>Processador Intel® Xeon® Gold 6230 (cache de 27,5 M, 2,10 GHz) Product Specifications.pdf</p>	<p>Página 1 do catálogo "frequência/Nº threads"</p>
<p>1.18.3 A solução entregue deverá aumentar progressivamente a quantidade de núcleos total fornecida de maneira proporcional à quantidade de recursos demandados</p>	<p>Pag. 9 (VMWARE VSAN)</p>	<p><b>Declaração técnica</b> Item: A solução entregue aumenta progressivamente a quantidade de núcleos total fornecida de maneira proporcional à</p>	<p>Página 01 da <b>Declaração técnica</b></p>



[www.lta-rh.com.br](http://www.lta-rh.com.br) | [comercial@lta-rh.com.br](mailto:comercial@lta-rh.com.br)

**Matriz** | Av. Ipiranga, 2640 | Santa Cecília | Porto Alegre | RS | Brasil | CEP 90610-000 | (51) 3382.7700/3094.1500  
**Filial DF** | SCN, Quadra 04 | Bloco B | Sala 1251/1252 | Distrito Federal | DF | Brasil | CEP:70714-900 | (61) 3034-3004  
**Filial SP** | Av. Paulista, 1636 | Conj.706 | Bela Vista | São Paulo | SP | Brasil | CEP 01310-200 | (11) 2391-9461  
**Filial RJ** | Praia de Botafogo 501 | Blc I Sala 101 | Botafogo | Rio de Janeiro | RJ | CEP 22.250-040 | (21) 2586-6000  
**Filial PR** | Rua Comendador Araújo 499 | CONJ 1007 | Centro | Curitiba | PR | CEP: 80.420-000 | (41) 99104-3240  
**Filial MG** | Av. Do Contorno, 6594 | 705 | Belo Horizonte | MG | CEP 30110-044 | (31) 3555-3477

pelos componentes relativos à infraestrutura interna de controle e gerenciamento da solução operando nos nós (exemplo: máquinas virtuais para administração do armazenamento SDS - Software-Defined Storage );		quantidade de recursos demandados pelos componentes relativos à infraestrutura interna de controle e gerenciamento da solução operando nos nós;	
1.18.4 Cada nó da solução deve possuir controladora de memória com suporte a DDR4 de, no mínimo, 2400MHz e 6 canais de memória;	Pag. 9 <b>INTEL® XEON® GOLD 6230 2.1G, 20C/40T, 10.4GT/S, 27.5M CACHE, TURBO, HT (125W) DDR4-2933</b>	Processador Intel® Xeon® Gold 6230 (cache de 27,5 M, 2,10 GHz) Product Specifications.pdf	Página 2 Velocidade máxima de memória / Nº máximo de canais de memória
1.19 A solução deverá operar com pentes de memória RAM tipo DDR4 RDIMM (Registered DIMM) ou LRDIMM (Load Reduced DIMM), com tecnologia de correção ECC (Error Correcting Code) e velocidade de, no mínimo, 2400MHz, e possuir as seguintes características:	Pag. 9 32GB RDIMM, 2933MT/S, DUAL RANK,BCC	Item: Replacing and adding hardware Subitem: System memory  <b>Declaração técnica</b> Item: Memória RAM com tecnologia de correção ECC (Error Correcting Code);	Página 29 do catálogo "docu87343_VxRail-Appliances-on-14th-Generation-PowerEdge-Servers-E-Series-Owner's-Manual"  Página 01 da <b>Declaração técnica</b>
1.19.1 A solução deverá entregar uma capacidade líquida mínima de memória RAM (retirando-se o montante utilizado para operação da solução) de 96 GB considerando inclusive cenário de operação parcial da solução (falha de algum nó);	Pag. 9 32GB RDIMM, 2933MT/S, DUAL RANK,BCC (total de 128GB por nó)	De acordo, atendemos ao solicitado, conforme consta nesta proposta	



[www.lta-rh.com.br](http://www.lta-rh.com.br) | [comercial@lta-rh.com.br](mailto:comercial@lta-rh.com.br)

**Matriz** | Av. Ipiranga, 2640 | Santa Cecília | Porto Alegre | RS | Brasil | CEP 90610-000 | (51) 3382.7700/3094.1500

**Filial DF** | SCN, Quadra 04 | Bloco B | Sala 1251/1252 | Distrito Federal | DF | Brasil | CEP:70714-900 | (61) 3034-3004

**Filial SP** | Av. Paulista, 1636 | Conj.706 | Bela Vista | São Paulo | SP | Brasil | CEP 01310-200 | (11) 2391-9461

**Filial RJ** | Praia de Botafogo 501 | Blc I Sala 101 | Botafogo | Rio de Janeiro | RJ | CEP 22.250-040 | (21) 2586-6000

**Filial PR** | Rua Comendador Araújo 499 | CONJ 1007 | Centro | Curitiba | PR | CEP: 80.420-000 | (41) 99104-3240

**Filial MG** | Av. Do Contorno, 6594 | 705 | Belo Horizonte | MG | CEP 30110-044 | (31) 3555-3477



<p>1.20 Os recursos de armazenamento deverão ser compartilhados entre todos os servidores e appliances da solução por meio de armazenamento definido por software - SDS;</p>	<p>Pag. 9 (VMWARE VSAN)</p>	<p>Item: VSAN subitem: vSAN aggregates locally attached disks of hosts in a vSphere cluster to create a pool of distributed shared storage</p>	<p>Página 57 do catálogo "h15104-vxrail-appliance-techbook"</p>
<p>1.21 A solução deverá entregar uma capacidade líquida de armazenamento de, no mínimo, 12 TB, considerando inclusive cenário de operação parcial da solução (falha de algum nó);</p>	<p>Pag. 9 3,84TB SSD qtd: 5 para cada appliance</p>	<p>Conforme proposta técnica comercial serão entregues 5 discos SSD de 3,84TB totalizando área líquida (após exclusão de 25% de slack space) de aproximadamente 12,825TB.</p>	<p>Conforme proposta técnica e comercial serão entregues 5 discos SSD de 3,84TB totalizando área líquida de 12,85TB. Para o cálculo consideramos o seguinte: Número de nós= 2 Quantidade de discos do group 1= 3 Quantidade de discos do group 2= 2 Número total de discos de capacidade=5 Tamanho do disco em TB (base 2)=3,42TB Total Bruto por node (base 2)=17,1TB Total Bruto no cluster de 2 nós (base 2)=34,2TB Área de "Slack Space" a ser descontada - 25% Percentual considerado para garantia de replicação de dados(FTT=1)-&gt;50%</p> <p><b>Total Líquido em TB disponível para VMs (base 2) = (ÁREA BRUTA TOTAL)*75%*50%=12,825TB</b></p>
<p>1.21.1 A tecnologia de discos de armazenamento da solução poderá ser compostas de qualquer tecnologia de discos, de forma híbrida ou não, desde que não seja totalmente composta somente por discos NL-SAS;</p>	<p>Não se aplica à nossa oferta</p>		



[www.lta-rh.com.br](http://www.lta-rh.com.br) | [comercial@lta-rh.com.br](mailto:comercial@lta-rh.com.br)

**Matriz** | Av. Ipiranga, 2640 | Santa Cecília | Porto Alegre | RS | Brasil | CEP 90610-000 | (51) 3382.7700/3094.1500

**Filial DF** | SCN, Quadra 04 | Bloco B | Sala 1251/1252 | Distrito Federal | DF | Brasil | CEP:70714-900 | (61) 3034-3004

**Filial SP** | Av. Paulista, 1636 | Conj.706 | Bela Vista | São Paulo | SP | Brasil | CEP 01310-200 | (11) 2391-9461

**Filial RJ** | Praia de Botafogo 501 | Blc I Sala 101 | Botafogo | Rio de Janeiro | RJ | CEP 22.250-040 | (21) 2586-6000

**Filial PR** | Rua Comendador Araújo 499 | CONJ 1007 | Centro | Curitiba | PR | CEP: 80.420-000 | (41) 99104-3240

**Filial MG** | Av. Do Contorno, 6594 | 705 | Belo Horizonte | MG | CEP 30110-044 | (31) 3555-3477

<p>1.21.2 No caso da composição dos discos dos nós da solução ser composta por discos NL-SAS, deve-se respeitar a proporção de, no mínimo, 30% de armazenamento líquido em discos SSD;</p>	<p>Não se aplica à nossa oferta</p>		
<p>1.22 Deverá ser permitida a troca de discos avariados, sem interrupção das operações de I/O das aplicações que acessarão os dados;</p>	<p>Pag. 9 (VXRAIL SOFTWARE)</p>	<p>VxRail Manager</p>	<p>Página 34 e 35 do catálogo "h15104-vxrail-appliance-techbook"</p>
<p>1.23 A falha isolada de um nó da solução não poderá impactar a disponibilidade da infraestrutura de armazenamento para as máquinas virtuais;</p>	<p>Pag. 9 (VMWARE VSAN)</p>	<p>Item: Storage Policy Based Management (SPBM) Subitem: Failures to tolerate (hosts)</p>	<p>Página 63 do catálogo "h15104-vxrail-appliance-techbook"</p>
<p>1.24 A solução deverá permitir a operação em grupo de, pelo menos, 2 appliances do mesmo tipo;</p>	<p>Pag. 9 (VMWARE VSAN)</p>	<p><b>Declaração técnica</b> (Item: A solução permite a operação em grupo de, pelo menos, 2 appliances do mesmo tipo;</p>	<p>Página 01 da <b>Declaração técnica</b></p>
<p>1.25 A solução deverá possuir console de administração web que deverá permitir integração com Microsoft Active Directory para autenticação;</p>	<p>Pag. 9 (VXRAIL SOFTWARE)</p>	<p>Item: Other administrative tasks Subitem: Active Directory user management tasks</p>	<p>Página 60 do catálogo "docu86660"</p>
<p>1.26 Deverá prover monitoração automática e periódica da solução, com o envio de notificações preventivamente em caso de falhas, notificando o suporte do fornecedor a tomar medidas preventivas e acordadas com o</p>	<p>Pag. 9 (VXRAIL SOFTWARE)</p>	<p>VxRail Manager ESRS details</p>	<p>Página 35 e 36 do catálogo "h15104-vxrail-appliance-techbook"</p>



[www.lta-rh.com.br](http://www.lta-rh.com.br) | [comercial@lta-rh.com.br](mailto:comercial@lta-rh.com.br)

**Matriz** | Av. Ipiranga, 2640 | Santa Cecília | Porto Alegre | RS | Brasil | CEP 90610-000 | (51) 3382.7700/3094.1500  
**Filial DF** | SCN, Quadra 04 | Bloco B | Sala 1251/1252 | Distrito Federal | DF | Brasil | CEP:70714-900 | (61) 3034-3004  
**Filial SP** | Av. Paulista, 1636 | Conj.706 | Bela Vista | São Paulo | SP | Brasil | CEP 01310-200 | (11) 2391-9461  
**Filial RJ** | Praia de Botafogo 501 | Blc I Sala 101 | Botafogo | Rio de Janeiro | RJ | CEP 22.250-040 | (21) 2586-6000  
**Filial PR** | Rua Comendador Araújo 499 | CONJ 1007 | Centro | Curitiba | PR | CEP: 80.420-000 | (41) 99104-3240  
**Filial MG** | Av. Do Contorno, 6594 | 705 | Belo Horizonte | MG | CEP 30110-044 | (31) 3555-3477

CONTRATANTE a fim de evitar tempo de inatividade e impactos na produção;			
1.27 Deverá oferecer portal de acesso do próprio fornecedor para download de atualizações e de softwares agregados à solução a fim de atender rapidamente demandas dos negócios;	Pag. 9 (VXRAIL E560)	Item: Customer upgradeable software <a href="https://support.emc.com/products/40127?siteLocale=en_US">https://support.emc.com/products/40127?siteLocale=en_US</a>	Página 36 e 37 do catálogo "h15104-vxrail-appliance-techbook"
1.28 A ferramenta de gerenciamento deverá fornecer um dashboard reportando a utilização dos recursos da solução como CPU, memória e armazenamento;	Pag. 9 (VXRAIL SOFTWARE)	Item: Metric Groups in vSphere Table 1-3. Metric Groups Item: Data Collection Intervals Table 1-4. Collection Intervals	Página 10 e 11 do catálogo "vsphere-esxi-vcserver-67-monitoring-performance-guide"
1.29 Comprovar que o equipamento não é fabricado utilizando as seguintes substâncias nocivas ao meio ambiente nas quantidades permitidas pela diretiva RoHS: chumbo (Pb), cádmio (Cd), mercúrio (Hg), cromo hexavalente (Hex-CR, bifenilos polibromados (PBBs) e éteres defenil-polibromados (PBDEs);	Pag. 9 (VXRAIL E560)	<b>Declaração técnica</b> Item: O equipamento não é fabricado utilizando as seguintes substâncias nocivas ao meio ambiente nas quantidades permitidas pela diretiva RoHS: chumbo (Pb), cádmio (Cd), mercúrio (Hg), cromo hexavalente (Hex-CR, bifenilos polibromados (PBBs) e éteres defenil-polibromados (PBDEs);	Página 01 da <b>Declaração técnica</b>
<b>ITEM 2 - INSTALAÇÃO DA SOLUÇÃO DE HIPERCONVERGÊNCIA</b>			



[www.lta-rh.com.br](http://www.lta-rh.com.br) | [comercial@lta-rh.com.br](mailto:comercial@lta-rh.com.br)

**Matriz** | Av. Ipiranga, 2640 | Santa Cecília | Porto Alegre | RS | Brasil | CEP 90610-000 | (51) 3382.7700/3094.1500

**Filial DF** | SCN, Quadra 04 | Bloco B | Sala 1251/1252 | Distrito Federal | DF | Brasil | CEP:70714-900 | (61) 3034-3004

**Filial SP** | Av. Paulista, 1636 | Conj.706 | Bela Vista | São Paulo | SP | Brasil | CEP 01310-200 | (11) 2391-9461

**Filial RJ** | Praia de Botafogo 501 | Blc I Sala 101 | Botafogo | Rio de Janeiro | RJ | CEP 22.250-040 | (21) 2586-6000

**Filial PR** | Rua Comendador Araújo 499 | CONJ 1007 | Centro | Curitiba | PR | CEP: 80.420-000 | (41) 99104-3240

**Filial MG** | Av. Do Contorno, 6594 | 705 | Belo Horizonte | MG | CEP 30110-044 | (31) 3555-3477

<p>2.1 A instalação da solução de hiperconvergência deverá ocorrer nos locais conforme distribuição no item 05 do Anexo I deste Edital I;</p>	<p>De acordo, atendemos ao solicitado, conforme consta nesta proposta</p>
<p>2.2 Todo ferramental necessário para execução dos serviços de instalação, configuração inicial, incluindo softwares, máquinas virtuais, equipamentos ou ferramentas, bem como eventuais materiais necessários para ligações temporárias, são de inteira responsabilidade da CONTRATADA;</p>	<p>De acordo, atendemos ao solicitado, conforme consta nesta proposta</p>



[www.lta-rh.com.br](http://www.lta-rh.com.br) | [comercial@lta-rh.com.br](mailto:comercial@lta-rh.com.br)

**Matriz** | Av. Ipiranga, 2640 | Santa Cecília | Porto Alegre | RS | Brasil | CEP 90610-000 | (51) 3382.7700/3094.1500

**Filial DF** | SCN, Quadra 04 | Bloco B | Sala 1251/1252 | Distrito Federal | DF | Brasil | CEP:70714-900 | (61) 3034-3004

**Filial SP** | Av. Paulista, 1636 | Conj.706 | Bela Vista | São Paulo | SP | Brasil | CEP 01310-200 | (11) 2391-9461

**Filial RJ** | Praia de Botafogo 501 | Blc I Sala 101 | Botafogo | Rio de Janeiro | RJ | CEP 22.250-040 | (21) 2586-6000

**Filial PR** | Rua Comendador Araújo 499 | CONJ 1007 | Centro | Curitiba | PR | CEP: 80.420-000 | (41) 99104-3240

**Filial MG** | Av. Do Contorno, 6594 | 705 | Belo Horizonte | MG | CEP 30110-044 | (31) 3555-3477

<p>2.3 A CONTRATANTE disponibilizará o espaço no CPD, refrigeração suficiente para comportar a solução a ser adquirida, assim como, a infraestrutura elétrica até a posição onde será instalado a solução, com capacidades (corrente e tensão) suficientes de suportar todos os equipamentos.</p>	<p>De acordo, atendemos ao solicitado, conforme consta nesta proposta</p>
<p>2.4 Entende-se por instalação, para efeito deste projeto: 2.4.1 Adequação das tomadas elétricas, caso necessário, aos requisitos da solução e características do ambiente, englobando fornecimento de materiais (cabearno elétrico, régua, adaptadores etc.) necessários ao perfeito funcionamento da solução ofertada;</p>	<p>De acordo, atendemos ao solicitado, conforme consta nesta proposta</p>



[www.lta-rh.com.br](http://www.lta-rh.com.br) | [comercial@lta-rh.com.br](mailto:comercial@lta-rh.com.br)

**Matriz** | Av. Ipiranga, 2640 | Santa Cecília | Porto Alegre | RS | Brasil | CEP 90610-000 | (51) 3382.7700/3094.1500

**Filial DF** | SCN, Quadra 04 | Bloco B | Sala 1251/1252 | Distrito Federal | DF | Brasil | CEP:70714-900 | (61) 3034-3004

**Filial SP** | Av. Paulista, 1636 | Conj.706 | Bela Vista | São Paulo | SP | Brasil | CEP 01310-200 | (11) 2391-9461

**Filial RJ** | Praia de Botafogo 501 | Blc I Sala 101 | Botafogo | Rio de Janeiro | RJ | CEP 22.250-040 | (21) 2586-6000

**Filial PR** | Rua Comendador Araújo 499 | CONJ 1007 | Centro | Curitiba | PR | CEP: 80.420-000 | (41) 99104-3240

**Filial MG** | Av. Do Contorno, 6594 | 705 | Belo Horizonte | MG | CEP 30110-044 | (31) 3555-3477



<p>2.4.2 Instalação física e lógica do equipamento, com fornecimento de cabos, UTP Cat. 6 para a totalidade das portas da solução e de gerenciamento, interligando os mesmos às redes LAN;</p>	<p>De acordo, atendemos ao solicitado, conforme consta nesta proposta</p>
<p>2.4.3 A metragem dos cabos fornecidos devem ser de, no mínimo, 5 metros;</p>	<p>De acordo, atendemos ao solicitado, conforme consta nesta proposta</p>
<p>2.4.4 As cores dos cabos fornecidos devem ser:</p>	<p>De acordo, atendemos ao solicitado, conforme consta nesta proposta</p>
<p>2.4.4.1 4 (quatro) cabos azuis;</p>	<p>De acordo, atendemos ao solicitado, conforme consta nesta proposta</p>
<p>2.4.4.2 1 (um) cabo cinza (a ser usado para gerenciamento do servidor);</p>	<p>De acordo, atendemos ao solicitado, conforme consta nesta proposta</p>
<p>2.5 A solução deverá ser instalada na última versão de firmware disponível pelo fabricante;</p>	<p>De acordo, atendemos ao solicitado, conforme consta nesta proposta</p>



[www.lta-rh.com.br](http://www.lta-rh.com.br) | [comercial@lta-rh.com.br](mailto:comercial@lta-rh.com.br)

**Matriz** | Av. Ipiranga, 2640 | Santa Cecília | Porto Alegre | RS | Brasil | CEP 90610-000 | (51) 3382.7700/3094.1500

**Filial DF** | SCN, Quadra 04 | Bloco B | Sala 1251/1252 | Distrito Federal | DF | Brasil | CEP:70714-900 | (61) 3034-3004

**Filial SP** | Av. Paulista, 1636 | Conj.706 | Bela Vista | São Paulo | SP | Brasil | CEP 01310-200 | (11) 2391-9461

**Filial RJ** | Praia de Botafogo 501 | Blc I Sala 101 | Botafogo | Rio de Janeiro | RJ | CEP 22.250-040 | (21) 2586-6000

**Filial PR** | Rua Comendador Araújo 499 | CONJ 1007 | Centro | Curitiba | PR | CEP: 80.420-000 | (41) 99104-3240

**Filial MG** | Av. Do Contorno, 6594 | 705 | Belo Horizonte | MG | CEP 30110-044 | (31) 3555-3477

<p>2.6 A instalação do equipamento deverá ser executada pela contratada, prestados por profissionais devidamente qualificados e/ou diretamente pelo fabricante dos equipamentos;</p>	<p>De acordo, atendemos ao solicitado, conforme consta nesta proposta</p>
<p>2.7 Entende-se por configuração inicial, para efeito deste projeto:</p>	
<p>2.7.1 Elaboração em conjunto com a equipe técnica do CONTRATANTE, de projeto de configuração, segundo as melhores práticas do fabricante e considerando as demandas e características dos serviços do CONTRATANTE;</p>	<p>De acordo, atendemos ao solicitado, conforme consta nesta proposta</p>
<p>2.7.2 Realização da configuração inicial da solução ofertada, segundo projeto, e conforme padrão de endereçamento IP a ser fornecido pelo CONTRATANTE;</p>	<p>De acordo, atendemos ao solicitado, conforme consta nesta proposta</p>



[www.lta-rh.com.br](http://www.lta-rh.com.br) | [comercial@lta-rh.com.br](mailto:comercial@lta-rh.com.br)

**Matriz** | Av. Ipiranga, 2640 | Santa Cecília | Porto Alegre | RS | Brasil | CEP 90610-000 | (51) 3382.7700/3094.1500

**Filial DF** | SCN, Quadra 04 | Bloco B | Sala 1251/1252 | Distrito Federal | DF | Brasil | CEP:70714-900 | (61) 3034-3004

**Filial SP** | Av. Paulista, 1636 | Conj.706 | Bela Vista | São Paulo | SP | Brasil | CEP 01310-200 | (11) 2391-9461

**Filial RJ** | Praia de Botafogo 501 | Blc I Sala 101 | Botafogo | Rio de Janeiro | RJ | CEP 22.250-040 | (21) 2586-6000

**Filial PR** | Rua Comendador Araújo 499 | CONJ 1007 | Centro | Curitiba | PR | CEP: 80.420-000 | (41) 99104-3240

**Filial MG** | Av. Do Contorno, 6594 | 705 | Belo Horizonte | MG | CEP 30110-044 | (31) 3555-3477

<p>2.7.3 Migração de uma máquina virtual existente no ambiente do CONTRATANTE para a solução ofertada e realização de teste de falha ao simular falha de um nó sem acarretar indisponibilidade do serviço provido por essa máquina virtual.</p>	<p>De acordo, atendemos ao solicitado, conforme consta nesta proposta</p>
<p>2.8 Fornecimento e configuração de software de SDS padrão do fabricante;</p>	<p>De acordo, atendemos ao solicitado, conforme consta nesta proposta</p>

Brasília, 1º de novembro de 2019.

\_\_\_\_\_  
ALEXANDER C. BARCELOS  
DIRETOR COMERCIAL  
CPF: 594.509.830-20 | RG: 2035263058



[www.lta-rh.com.br](http://www.lta-rh.com.br) | [comercial@lta-rh.com.br](mailto:comercial@lta-rh.com.br)

**Matriz** | Av. Ipiranga, 2640 | Santa Cecília | Porto Alegre | RS | Brasil | CEP 90610-000 | (51) 3382.7700/3094.1500

**Filial DF** | SCN, Quadra 04 | Bloco B | Sala 1251/1252 | Distrito Federal | DF | Brasil | CEP:70714-900 | (61) 3034-3004

**Filial SP** | Av. Paulista, 1636 | Conj.706 | Bela Vista | São Paulo | SP | Brasil | CEP 01310-200 | (11) 2391-9461

**Filial RJ** | Praia de Botafogo 501 | Blc I Sala 101 | Botafogo | Rio de Janeiro | RJ | CEP 22.250-040 | (21) 2586-6000

**Filial PR** | Rua Comendador Araújo 499 | CONJ 1007 | Centro | Curitiba | PR | CEP: 80.420-000 | (41) 99104-3240

**Filial MG** | Av. Do Contorno, 6594 | 705 | Belo Horizonte | MG | CEP 30110-044 | (31) 3555-3477

**ANEXO IV - PREGÃO ELETRÔNICO SRP Nº 40/2019  
DECLARAÇÃO DE VISTORIA TÉCNICA**

**Ao**  
**TRIBUNAL REGIONAL FEDERAL DA PRIMEIRA REGIÃO – TRF 1ª REGIÃO**  
**Ref.: PREGÃO ELETRÔNICO RP Nº 40/2019**  
**Proposta N.º 281/19**

Declaro, para fins de participação na licitação em epígrafe, que vistoriei o ambiente tecnológico do CONTRATANTE onde serão entregues e instalados os equipamentos, bem com onde serão prestados os serviços assistência técnica e treinamento, tomei conhecimento das informações necessárias à execução do contrato.

Por fim proclamo estar ciente da complexidade dos serviços, dos termos e condições descritos no respectivo edital e seus anexos.

Declaro que todas as dúvidas que porventura foram por mim questionadas foram respondidas pela equipe técnica do TRF1 e que marquei de próprio punho os itens abaixo.

	Realizei vistoria no ambiente e condições das instalações.
	Tomei conhecimento dos procedimentos adotados, documentação existente, modelos de acompanhamento, recomendações e normatizações da JF1.
	Estou ciente do grau de dificuldade e a devida especialização técnica necessária para a execução dos serviços a serem contratados.
	Estou ciente do grau de dificuldade e a devida especialização técnica necessária para a execução dos serviços a serem contratados.

x	Declaro que renuncio à Visita Técnica dos ambientes tecnológicos do CONTRATANTE em que serão entregues os equipamentos e prestados os serviços constantes do objeto do Edital nº 40/2019, e que tenho pleno conhecimento das condições e peculiaridades inerentes à natureza dos trabalhos, assumindo total responsabilidade pela ocorrência de eventuais prejuízos em virtude de minha omissão na verificação dos ambientes tecnológicos do CONTRATANTE, com vistas a proteger o interesse da Administração, nos termos do Acórdão 1174/2008 – Plenário - TCU.
---	---

Brasília, 1º de novembro de 2019.

\_\_\_\_\_  
ALEXANDER C. BARCELOS  
DIRETOR COMERCIAL  
CPF: 594.509.830-20 | RG: 2035263058



[www.lta-rh.com.br](http://www.lta-rh.com.br) | [comercial@lta-rh.com.br](mailto:comercial@lta-rh.com.br)

**Matriz** | Av. Ipiranga, 2640 | Santa Cecília | Porto Alegre | RS | Brasil | CEP 90610-000 | (51) 3382.7700/3094.1500

**Filial DF** | SCN, Quadra 04 | Bloco B | Sala 1251/1252 | Distrito Federal | DF | Brasil | CEP:70714-900 | (61) 3034-3004

**Filial SP** | Av. Paulista, 1636 | Conj.706 | Bela Vista | São Paulo | SP | Brasil | CEP 01310-200 | (11) 2391-9461

**Filial RJ** | Praia de Botafogo 501 | Blc I Sala 101 | Botafogo | Rio de Janeiro | RJ | CEP 22.250-040 | (21) 2586-6000

**Filial PR** | Rua Comendador Araújo 499 | CONJ 1007 | Centro | Curitiba | PR | CEP: 80.420-000 | (41) 99104-3240

**Filial MG** | Av. Do Contorno, 6594 | 705 | Belo Horizonte | MG | CEP 30110-044 | (31) 3555-3477

DECLARAÇÃO

Ao  
TRIBUNAL REGIONAL FEDERAL DA PRIMEIRA REGIÃO – TRF 1ª REGIÃO  
Ref.: PREGÃO ELETRÔNICO RP Nº 40/2019  
Proposta N.º 281/19

A LTA-RH Informática Comércio, Representações Ltda., com sede na Av. Ipiranga, 2640, Santa Cecília, Porto Alegre/RS, CEP: 90610-000, inscrita sob o CNPJ nº 94.316.916/0001-07 e com filial em SCN, quadra 04, bloco B, Sala 1251/1252, Distrito Federal/DF, CEP: 70714-900, inscrita sob o CNPJ nº 94.316.916/0005-22, sob pena da lei, declara o que segue:

A **LTA-RH INFORMÁTICA COMÉRCIO, REPRESENTAÇÕES LTDA.**, está autorizada a comercializar e prestar assistência técnica no período de garantia nos equipamentos propostos, conforme as exigências e prazos contidos no edital e seus anexos, com Centro de Assistência Técnica em São Paulo, com atendimento em todo estado, bem como nos locais de entrega e instalação, por sua Central de Atendimento de Tele - Suporte com discagem gratuita pelo telefone 0800-5105820, através do site: <https://app01.lta-rh.com.br/fusion/portal> ou através do responsável técnico Sr. Daniel Santos no e-mail: [daniel\\_santos@lta-rh.com.br](mailto:daniel_santos@lta-rh.com.br).

- Que atende a todos os requisitos do edital e seus anexos.
- Garantia de 60 (sessenta);
- Prazo de atendimento técnico:
  - 02 (dois) dias úteis para equipamentos instalados no TRF1 e seções judiciais;
  - 03 (três) dias úteis para equipamentos instalados nas subseções judiciais;
  - 05 (cinco) dias corridos especificamente para equipamentos instalados nas subseções judiciais a seguir: Oiapoque (PA), Laranjal do Jari (AP), Tabatinga (AM) e Tefé (AM).
- Abertura de 3 (três) chamados de assistência técnica no período de 20 (vinte) dias corridos por causas similares ou defeitos de mesma natureza.

Brasília, 1º de novembro de 2019.

---

ALEXANDER C. BARCELOS  
DIRETOR COMERCIAL  
CPF: 594.509.830-20 | RG: 2035263058



[www.lta-rh.com.br](http://www.lta-rh.com.br) | [comercial@lta-rh.com.br](mailto:comercial@lta-rh.com.br)

**Matriz** | Av. Ipiranga, 2640 | Santa Cecília | Porto Alegre | RS | Brasil | CEP 90610-000 | (51) 3382.7700/3094.1500

**Filial DF** | SCN, Quadra 04 | Bloco B | Sala 1251/1252 | Distrito Federal | DF | Brasil | CEP:70714-900 | (61) 3034-3004

**Filial SP** | Av. Paulista, 1636 | Conj.706 | Bela Vista | São Paulo | SP | Brasil | CEP 01310-200 | (11) 2391-9461

**Filial RJ** | Praia de Botafogo 501 | Blc I Sala 101 | Botafogo | Rio de Janeiro | RJ | CEP 22.250-040 | (21) 2586-6000

**Filial PR** | Rua Comendador Araújo 499 | CONJ 1007 | Centro | Curitiba | PR | CEP: 80.420-000 | (41) 99104-3240

**Filial MG** | Av. Do Contorno, 6594 | 705 | Belo Horizonte | MG | CEP 30110-044 | (31) 3555-3477





### Declaração do Fabricante

**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 **LTA-RH INFORMÁTICA COMÉRCIO, REPRESENTAÇÕES LTDA**, no certame licitatório n. Pregão Eletrônico RP nº 40/2019, promovido pelo TRIBUNAL REGIONAL FEDERAL DA PRIMEIRA REGIÃO, vem, através desta, declarar que:

- o(s) modelo(s) **VXRAIL E560**, possui(em) garantia de 60 meses, on-site, com atendimento telefônico 24 horas por dia, 7 dias na semana.

Declaramos, ainda, que:

- **O SLA (tempos de atendimento e/ou solução) exigido no edital será de responsabilidade exclusiva da LTA-RH INFORMÁTICA COMÉRCIO, REPRESENTAÇÕES LTDA, por sua central de atendimento de tele-suporte com discagem pelo telefone 0800-5105820 ou (51) 3382-7700.**

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

- A **LTA-RH INFORMÁTICA COMÉRCIO, REPRESENTAÇÕES LTDA** está autorizada a comercializar os equipamentos propostos para esse certame.

Eldorado do Sul, 10 de outubro de 2019

**Dell Computadores do Brasil Ltda**

**Geoquisa Idalina Rodrigues Terra – Gerente de Vendas**

**DELL Computadores do Brasil Ltda.**

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



Eldorado do Sul, 10 de outubro de 2019

À  
LTA-RH INFORMÁTICA, COMÉRCIO E REPRESENTAÇÕES LTDA  
A/C Sr. FABIANO COSTA BARCELOS

Ref: TRIBUNAL REGIONAL FEDERAL DA 1 REGIAO - TRF 1 REGIAO PREGÃO ELETRÔNICO RP  
Nº 40/2019

### DECLARAÇÃO TÉCNICA

**DELL COMPUTADORES DO BRASIL LTDA. ("Dell")**, inscrita no CNPJ/MF sob o nº 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 EMC VXRAIL E560**

Solução de hiperconvergência baseada em appliances, segundo a convenção da Associação da Indústria de Redes de Armazenamento – SNIA (Storage Networking Industry Association);

Possui botão liga/desliga com proteção para prevenir o desligamento acidental;

Será entregue junto com cada módulo kit de fixação no rack, do tipo retrátil, permitindo o deslizamento do servidor a fim de facilitar sua manutenção;

Possui Ventilação hotpluggable;

Cada fontes têm potência suficiente para atender adequadamente a configuração máxima do equipamento entregue, acompanhado de cabo de alimentação.

O equipamento não é fabricado utilizando as seguintes substâncias nocivas ao meio ambiente nas quantidades permitidas pela diretiva RoHS: chumbo (Pb), cádmio (Cd), mercúrio (Hg), cromo hexavalente (Hex-CR), bifenilos polibromados (PBBs) e éteres defenil-polibromados (PBDEs);

Memória RAM com tecnologia de correção ECC (Error Correcting Code);

A solução está estruturada de forma a suportar a implementação de ambiente de virtualização em alta disponibilidade, conforme as boas práticas do fabricante da solução de virtualização;

Toda a solução possui suporte centralizado em uma única central de atendimento do fornecedor da solução de hardware;

A solução entregue aumenta progressivamente a quantidade de núcleos total fornecida de maneira proporcional à quantidade de recursos demandados pelos componentes relativos à infraestrutura interna de controle e gerenciamento da solução operando nos nós;

A solução permite a operação em grupo de, pelo menos, 2 appliances do mesmo tipo;

Atenciosamente,

**Dell Computadores do Brasil Ltda**

**Geoquiza Idalina Rodrigues Terra – Gerente de Vendas**

**DELL Computadores do Brasil Ltda.**

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

Produto:	Unidade de processamento digital de média capacidade, baseada em microprocessadores
Processo MCT/Data:	01200.004347/2006-31 de 21/08/2006
CNPJ da Incentivada:	72381189000625
Portaria MCT/MDIC/MF:	985, de 22/12/2006 DOU 26/12/2006 ()
Modelos:	POWEREDGE.(11-07-08: SERVIDOR POWEREDGE R 900)(03-05-10: POWEREDGE R910)(28-05-10: POWEREDGE M910)(28-10-11: KACE K1100) (28-10-11:KACE K1200) (28-10-11: KACE K2100) (28-10-11: KACE K2200) (06-12-11: KACE K1100 ADV) (06-12-11: KACE K2100 ADV) (09-02-12: POWEREDGE R620) (09-02-12: POWEREDGE T620) (09-02-12: POWEREDGE R720) (09-02-12: POWEREDGE M620) (09-02-12: POWEREDGE M520) (17-02-12: SERVIDOR POWEREDGE R720 XD) (17-02-12: SERVIDOR POWEREDGE R720 OEM) (17-02-12: SERVIDOR POWEREDGE R620 OEM) (11-06-12: POWEREDGE R420) (11-06-12: POWEREDGE T420) (11-06-12: POWEREDGE R520) (11-06-12: POWEREDGE R520 OEM) (25-06-12: SERVIDOR POWEREDGE M820) (30-08-12: KACE K1100S) (30-08-12: KACE K1100S ADV) (30-08-12: KACE K1200S) (30-08-12: KACE K2100S ADV) (30-08-12: KACE K2200S) (30-08-12: KACE 2100S) (09-04-13: KACE K3100S) (09-04-13: KACE K3100S ADV) (09-04-13: KACE K3200S) (14-06-13: DESKTOP OPTIPLEX 9020)(22-11-13: POWEREDGE FX2) (22-11-13: SERVIDOR BLADE FC620) (22-11-13: POWEREDGE FC620) (30-12-13: SERVIDOR POWEREDGE R820) (11-03-14: POWEREDGE R920) (16-03-14: SERVIDOR POWEREDGE R220) (10-07-14: SERVIDOR POWEREDGE T630) (10-07-14: SERVIDOR POWEREDGE R730) (10-07-14: SERVIDOR POWEREDGE R630) (10-09-14: SERVIDOR POWEREDGE R430) (10-09-14: SERVIDOR POWEREDGE R530) (10-09-14: SERVIDOR POWEREDGE T430) (10-09-14: SERVIDOR POWEREDGE M630) (10-09-14: SERVIDOR POWEREDGE FC630) (01-04-15: POWEREDGE R930) (17-08-15: POWEREDGE T130) (17-08-15: POWEREDGE T330) (17-08-15: POWEREDGE R230)(27-11-15: POWEREDGE R730XD) (18-03-16: XC630) (18-03-16: XC730XD) (18-03-16: POWEREDGE R830)(27-04-17: POWEREDGE R640) (27-04-17: POWEREDGE R740) (27-04-17: POWEREDGE R740XD) (27-04-17: POWEREDGE R940) (27-04-17: EMC VXRAIL E460) (27-04-17: EMC VXRAIL P470) (27-04-17: EMC VXRAIL S470) (27-04-17: EMC VXRAIL V470) (27-04-17: EMC VXRAIL E460F) (27-04-17: EMC VXRAIL P470F) (27-04-17: EMC VXRAIL V470F) (14-06-17: POWEREDGE T440) (14-06-17: POWEREDGE R440) (14-06-17: POWEREDGE FC640) (14-06-17: POWEREDGE M640) (14-06-17: POWEREDGE R540) (14-06-17: POWEREDGE T640) (15-09-17: XC640) (15-09-17: XC740XD) (02-10-17: EMC VXRAIL E560) (02-10-17: EMC VXRAIL E560F) (02-10-17: EMC VXRAIL P570) (02-10-17: EMC VXRAIL P570F) (02-10-17: EMC VXRAIL S570) (02-10-17: EMC VXRAIL V570) (02-10-17: EMC VXRAIL V570F) (23-03-18: POWEREDGE R840) (23-03-18: POWEREDGE R840) (06-04-18: POWEREDGE MX740C) (31-07-18: POWEREDGE T340) (31-07-18: POWEREDGE R240) (31-07-18: POWEREDGE T140) (23-04-19: POWEREDGE C6400) (23-04-19: POWEREDGE C6420)

Produto:	Unidade de processamento digital de pequena capacidade, baseada em microprocessadores
Processo MCT/Data:	01200.004347/2006-31 de 21/08/2006
CNPJ da Incentivada:	72381189000625
Portaria MCT/MDIC/MF:	985, de 22/12/2006 DOU 26/12/2006 ()

# Dell EMC VxRail™ Appliance

Version 4.5.x

## Administration Guide

REV 03

Copyright © 2016-2018 Dell EMC. All rights reserved.

Published February 2018

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

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS-IS." DELL 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 DELL SOFTWARE DESCRIBED IN THIS PUBLICATION REQUIRES AN APPLICABLE SOFTWARE LICENSE.

Dell, EMC, and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be the property of their respective owners.  
Published in the USA.

Dell EMC  
Hopkinton, Massachusetts 01748-9103  
1-508-455-1000 (in North America) 866-464-7381  
[www.DellEMC.com](http://www.DellEMC.com)

# CONTENTS

<b>Chapter 1</b>	<b>Introduction</b>	<b>5</b>
	Introduction.....	6
	Support.....	6
	Registering for online support.....	6
	Where to go for support resources.....	7
	Using the SolVe Desktop application for VxRail Series procedures....	7
<b>Chapter 2</b>	<b>Architecture overview</b>	<b>9</b>
	Overview of the VxRail Series.....	10
	The VxRail Appliance.....	10
	Storage and virtualization.....	11
	Features.....	11
	VxRail Appliance cluster expansion.....	12
	Using a VxRail Appliance with an external VMware vCenter server.....	13
<b>Chapter 3</b>	<b>Network planning</b>	<b>15</b>
	Pre-Engagement Questionnaire (PEQ).....	16
	10GbE or 1GbE switch networking.....	16
	Network requirements.....	17
	VxRail Appliance Network Configuration Table.....	17
<b>Chapter 4</b>	<b>Licensing your VxRail Appliance</b>	<b>21</b>
	Licensing overview.....	22
	Locating your partner activation code (PAC).....	22
	Logging into the activation portal.....	22
	Preparing your vSphere license for your VxRail Appliance.....	23
	Assigning a license to your VxRail Appliance.....	24
<b>Chapter 5</b>	<b>Working with VxRail Manager</b>	<b>25</b>
	Viewing the VxRail Manager dashboard.....	26
	Support resources.....	27
	Viewing Secure Remote Support (ESRS) information.....	28
	Using Support eServices.....	28
	Using the support knowledge base.....	29
	Viewing system events.....	29
	Monitoring logical system health.....	30
	Viewing physical system health.....	32
	Viewing and locating hardware events.....	33
	Viewing physical disk status.....	34
	Viewing physical node status.....	36
	Viewing power supply status.....	39
	Viewing physical NIC status.....	40
	Viewing VxRail Appliance system configuration and applications.....	41
	Config General tab.....	42
	Config System tab.....	52
	Using the VxRail Market.....	54



<b>Chapter 6</b>	<b>Other administrative tasks</b>	<b>59</b>
	Active Directory user management tasks.....	60
	Cluster management tasks.....	60
	Data center management tasks.....	60
	Datastore management tasks.....	61
	DNS management tasks.....	61
	ESRS management tasks.....	62
	ESXi host management tasks.....	62
	External vCSA and PSC management tasks.....	62
	Internal vCSA and PSC management tasks.....	63
	Migrating from a VxRail vCenter Server to a customer-supplied vCenter Server.....	64
	iDRAC configuration management tasks.....	65
	iSCSI management tasks.....	65
	Log server management tasks.....	66
	NFS management tasks.....	66
	NTP server management tasks.....	66
	vSphere Distributed Switch (VDS) management tasks.....	66
	VM folder management tasks.....	68
	vSphere Standard Switch (VSS) management tasks.....	68
	VxRail Manager application management tasks.....	69
<b>Chapter 7</b>	<b>External storage</b>	<b>71</b>
	NFS and VMFS over iSCSI datastores.....	72
	Setting up external NFS storage for a VxRail Appliance.....	72
	Setting up external VMFS over iSCSI storage for a VxRail Appliance.....	73
	Set up the two iSCSI network portgroups.....	73
	Create two VMkernel adapters on the hosts in the cluster.....	75
	Set up an active software iSCSI adapter.....	75
	Assign volumes on iSCSI storage.....	76
	Configure multipath for the iSCSI volumes.....	77
	Mount a new VMFS datastore over iSCSI.....	77
<b>Chapter 8</b>	<b>Replacing and adding hardware</b>	<b>79</b>

# CHAPTER 1

## Introduction

- Introduction..... 6
- Support..... 6

## Introduction

This document describes the VxRail™ Appliance, how it works, and how to perform administrative tasks.

The target audience for this document includes customers, field personnel, and partners who want to manage and operate a VxRail™ Appliance. This document is designed for people familiar with:

- Dell EMC systems and software
- VMware virtualization products
- Data center appliances and infrastructure

## Support

Create an Online Support account to get access to support and product resources for your VxRail Appliance.

If you already have an account, register your VxRail Appliance to access the available resources.

For convenience, you can link your Online Support account with VxRail Manager and access support resources without having to log in separately.

---

### Note

If you plan to set up ESRS, your Online Support account must be linked to VxRail Manager under the same party ID or the deployment will fail. Your appliance must also be in an installed state in the install Base.

---

## Registering for online support

Create an Online Support account to access support resources.

After you register, you can:

- Register your system
- Obtain product license files and software updates
- Download VxRail Series product documentation
- Download the Solve Desktop application for hardware replacement and upgrade procedures
- Browse the VxRail Series community and support information
- Link your support account for access to resources from within VxRail Manager

### Procedure

1. Point your Web browser to [emc.com/vxrailsupport](http://emc.com/vxrailsupport) (or [support.emc.com](http://support.emc.com)).
2. Click **Register here**.
3. Fill in the required information.

Support will send you a confirmation email, typically within 48 hours.

## Where to go for support resources

Access support resources for your VxRail Appliance by doing any of the following:

- Click the VxRail Manager **Support** tab.
- Point your Web browser to [emc.com/vxrailsupport](https://emc.com/vxrailsupport) (or [support.emc.com](https://support.emc.com)).

## Using the Solve Desktop application for VxRail Series procedures

Step-by-step instructions for procedures such as replacing hardware or performing system administrative tasks are available through the Solve Desktop application.

### **⚠ CAUTION**

**To avoid potential data loss, refer to the VxRail Series procedures in the Solve Desktop application before performing any hardware replacement or upgrade activity.**

---

To download the Solve Desktop application, go to <https://support.emc.com> and click **Solve** on the main page. Download and install the Solve Desktop application on your computer.

You must have an online support account to use the Solve Desktop application.

## Introduction

# CHAPTER 2

## Architecture overview

- Overview of the VxRail Series..... 10
- The VxRail Appliance..... 10
- Storage and virtualization..... 11
- Features..... 11
- VxRail Appliance cluster expansion..... 12
- Using a VxRail Appliance with an external VMware vCenter server..... 13



## Overview of the VxRail Series

The VxRail Series delivers virtualization, compute, and storage in a scalable, easy to manage, hyper-converged infrastructure appliance.

Your VxRail Appliance is built on Intel Xeon processor-based hardware with the VxRail Appliance software bundle, and support for other value-added software from Dell EMC and VMware.

The VxRail Appliance software bundle includes the following:

- VxRail™ Manager for deployment, configuration, and management
- VMware vSphere®, including ESXi
- VMware vCenter Server™
- VMware vSAN™ for storage
- VMware vRealize Log Insight™

Your VxRail Appliance also includes complementary applications such as CloudArray and RecoverPoint for Virtual Machines.

## The VxRail Appliance

The VxRail Appliance includes the appliance hardware, VxRail Manager, EMC Secure Remote Services (ESRS), and access to qualified software products.

### Appliance hardware

The VxRail Appliance consists of one to four nodes in a rackmount chassis. Each node has its own compute and storage resources.

For a list of available VxRail Appliance models refer to the VxRail website.

One or more network switches (10GbE or 1GbE depending on model), appropriate cables, and a workstation/laptop for the user interface are also required for use with your VxRail Appliance.

### VxRail Manager

VxRail Manager provides a software stack for software-defined data center (SDDC) building blocks including compute, network, storage, and management. VxRail Manager streamlines deployment, configuration, and management for easier initial setup and ongoing operations.

VxRail Manager also provides integration for services and support to help you get the most out of your VxRail Appliance.

With VxRail Manager you can:

- Monitor system health with deep hardware intelligence and graphical representation
- View appliance software versions and updates
- Access Online Support and eServices
- Access community resources such as the user forum and knowledgebase
- Use the VxRail™ Market to access qualified software products
- Perform maintenance operations such as replacing hardware, adding drives, and cycling power to the cluster or nodes

- Perform system software upgrades
- Expand the nodes in the cluster

## Storage and virtualization

The VxRail Series integrates VMware vSAN and virtualization to provide a complete hyper-converged infrastructure.

### Storage

VMware vSAN is integrated in your VxRail Appliance to provide Software-Defined Storage (SDS). vSAN is not a VSA, but is embedded in the ESXi hypervisor kernel's I/O data path. As a result, vSAN can deliver the higher performance with minimum CPU and memory overhead.

vSAN pools the VxRail Appliance's internal SSDs and HDDs on the ESXi hosts to present a single datastore for all hosts in the cluster. vSAN uses a highly available, distributed, object-based architecture. vSAN mirrors and distributes the individual virtual disk (VMDK) across the datastore.

### Virtualization

The VxRail Series lets virtualization infrastructure administrators manage storage on a per-VM basis. This VM-centric approach allows for storage policies to be defined at VM-level granularity for provisioning and load balancing. vSAN is fully integrated with vSphere, which simplifies setting up the availability, capacity, and performance policies.

For scale-out, VxRail Manager uses VMware Loudmouth auto-discovery capabilities, based on the RFC-recognized "Zero Network Configuration" protocol, to automatically discover and configure appliances on your network. Loudmouth runs on each ESXi host and in the VxRail Manager virtual machine. Loudmouth allows VxRail Manager to discover all the nodes and automate the configuration. Loudmouth requires IPv6 multicast. The IPv6 multicast communication is strictly limited to the management VLAN that the nodes use for communication.

## Features

The VxRail Series offers advanced features including automatic deployment, automatic scale out, fault tolerance, and diagnostic logging.

### Automatic deployment

The VxRail Manager fully automates the installation and configuration of all nodes in an appliance after you input the basic IP address information.

### Automatic scale-out

The VxRail Series provides automated scale-out functionality by detecting a new VxRail Appliance on the network. When a new VxRail Appliance is powered on you can add it to your existing cluster or create a new cluster, replicating the configuration, and expanding the datastore in a cluster.

New in VxRail Manager 4.5, scale-out for existing VxRail Appliances is now even faster with multi-node expansion. This feature allows multiple nodes to be added into the same cluster in parallel, reducing the time required to expand your cluster.

### Node failure tolerance

The VxRail Series tolerates node failures as defined by the vSAN policy. Refer to VMware vSAN documentation on FTT and FTM support for more information.

The VxRail Series implements the standard vSAN policy of one failure by default:

- An entire node can fail and the system will continue to function.
- Disk failure cannot affect more than one node.
- One cache disk can affect as many as six capacity disks (HDD/SSD).
- One network port on any node can fail without affecting the node.

Network failover is through the virtual switch configuration in ESXi. This is automatically configured by VxRail Manager during initial setup.

#### Logging and log bundles

The VxRail Series provides logging and log bundles through VxRail Manager. These logs provide operation and event information about VxRail Manager.

## VxRail Appliance cluster expansion

Your VxRail Appliance cluster can be scaled from a minimum of three nodes up to a maximum of 32 nodes. Additional nodes may be supported through the RPQ process (see below). The VxRail Manager automated installation and scale-out features make it easy to expand your cluster as your business demands grow.

With the VxRail Manager version 4.5 multi-node expansion feature, you can perform cluster expansion with up to six nodes in parallel, reducing the time required to achieve greater compute and storage capacity.

VxRail Appliance models hold from one to four nodes in a rackmount chassis. You may have a partially populated appliance chassis in the cluster. You can use the empty slots in the chassis for future expansion.

You can mix different VxRail Appliance models in the same cluster. You must adhere to the following guidelines when deploying a mixed cluster:

- The first four nodes in a cluster must be of the same type and configuration. (If there are only three nodes in the cluster they must all be the same type and configuration.)
- All appliances in the cluster must be running VxRail Manager version 4.5.0 or higher.
- First-generation appliances (sold under the VSPEX BLUE name) are not supported for VxRail Manager version 4.5.0 or higher.
- Appliances using 1GbE networking cannot be used in clusters with 10GbE networking.
- Hybrid nodes cannot be used in clusters with all-flash nodes.

---

#### Note

VxRail Appliance cluster scalability is supported to a maximum of 32 nodes. However, scalability to 64 nodes per cluster may be allowed. You must submit a request for product qualification (RPQ) for clusters over 32 nodes.

- VxRail Appliance G Series models support up to eight appliances. Up to 16 appliances may be supported through the RPQ process.
  - 1Gbs clusters support a maximum of eight nodes and two appliances.
- 

Work with your sales representative or reseller when planning your cluster expansion.

## Using a VxRail Appliance with an external VMware vCenter server

Your VxRail Appliance can join an existing external (customer-supplied) vCenter server during its initial configuration. This allows you to use a remote central vCenter server to manage multiple VxRail Appliances from a single pane of glass.

The external vCenter server can be:

- Physical or virtual
- Windows or VCSA
- Embedded, or non-embedded deployed with an external PSC

To join an existing external vCenter server, provide an existing datacenter and a non-conflicting cluster name during the initial configuration of the appliance. The VxRail Appliance joins the datacenter as a new vSAN cluster with the specified cluster name.

---

### Note

The following restrictions apply when using your VxRail Appliance with an external vCenter server:

- The external vCenter server can be VMware VCSA 6.5 or later, or VCSA 6.0 U2 or later.
  - Your VxRail Appliance cannot join an external vCenter server on another VxRail Appliance cluster. The vCenter server on one VxRail Appliance cluster cannot be used to manage other VxRail Appliance clusters.
  - You must provide an appropriate vCenter server license.
-

## Architecture overview

# CHAPTER 3

## Network planning

Before your VxRail Appliance is installed, your network switches must be properly configured. Consult with your sales representative or partner to prepare your switches and network before installation.

This chapter discusses the switch and network parameters you must consider to make sure that you properly setup up your switch, configure VLANs, reserve IP addresses, and so on.

- Pre-Engagement Questionnaire (PEG)..... 16
- 10GbE or 1GbE switch networking..... 16
- Network requirements..... 17
- VxRail Appliance Network Configuration Table..... 17



## Pre-Engagement Questionnaire (PEQ)

Before your appliance is installed, work with your pre-sales systems engineer, project manager, or partner to complete the VxRail Appliance Pre-Engagement Questionnaire (PEQ).

The PEQ tool is used to document the Professional Services project installation parameters needed to successfully install your appliance and to facilitate the communication between the different teams throughout the engagement workflow (pre-sales to project management to delivery).

## 10GbE or 1GbE switch networking

Your VxRail Appliance relies on your 10GbE or 1GbE switches for all of the networking between nodes in the cluster, and between the cluster and the rest of your infrastructure. It is crucial that your switches are configured properly for the VxRail Appliance to work.

Consult with your sales representative or partner to ensure your switch is properly set up to work with your VxRail Appliance according to the instructions provided by your switch manufacturer.

The following table lists the type of switches needed for different VxRail Appliance models.

**Table 1** VxRail Appliance network switch requirements

VxRail Appliance model	Network switch requirement
VxRail 60, VxRail 60F VxRail E460 VxRail G410, VxRail G410F VxRail P470 VxRail S470 VxRail V470	10GbE or 1GbE
VxRail 120, VxRail 120F VxRail 160, VxRail 160F VxRail 200, VxRail 200F VxRail 240F VxRail 280F VxRail E460F VxRail P470F VxRail V470F	10GbE

## Network requirements

Your sales representative or partner will work closely with you to ensure you have the following network resources available before installing the VxRail Appliance.

### 10GbE or 1GbE switches

You must have one or more 10GbE or 1GbE switches:

- Two or more network switches are recommended for failover, and for installations with more than four appliances.
- Two switch ports are required for each node (the VxRail 60 requires four ports per node). A fully-populated appliance can require eight switch ports.

### Cabling

If you have RJ45 NIC ports on your appliance you need:

- Two CAT6 or higher cables per node (these cables are included with RJ45-equipped appliances). A fully-populated appliance requires up to eight cables (the VxRail 60 requires 16 cables).

If you have SFP+ NIC ports on your appliance you need:

- Two compatible Twinax Direct-Attach-Copper (DAC) cables per node (these cables are not included; you must supply your own). Active DAC cables should be used if the cable is over 5 meters long.

or

- Two compatible fiber cables with appropriate transceivers per node (these cables and transceivers are not included; you must supply your own).
- A fully-populated appliance requires up to eight cables.

### Servers

You must have a DNS server for network address resolution and ESRS support.

## VxRail Appliance Network Configuration Table

Use the Network Configuration Table when planning and configuring your 10GbE or 1GbE switches for use with your VxRail Appliance.

This table lists the network parameters that are configured during VxRail Appliance setup. Use the table before setup to plan your configuration. Work with your sales representative or partner to use the table during setup to keep track of your entries.

Table 2 Network Configuration Table

Element	Category	Description	Example	Customer value
VxRail Appliance	Management VLAN ID (optionally modify)	Set a management VLAN on ESXi before you configure VxRail Appliance, otherwise management traffic will be untagged on the switch's Native VLAN	Native VLAN	
	VxRail Appliance Initial IP Address (optionally modify)	If you cannot reach the default VxRail Appliance initial IP address	192.168.10.200	

Table 2 Network Configuration Table (continued)

Element	Category	Description	Example	Customer value	
		(192.168.10.200/24), set an alternate IP address			
System	Global settings	Time zone			
		NTP server(s)			
		DNS server(s)			
	Active Directory (optional)	Domain			
		Username			
		Password			
	HTTP Proxy Settings (optional)	IP Address			
		Port			
		Username			
		Password			
Management	Hostnames	ESXi hostname prefix	host		
		Separator	None		
		Iterator	0X		
		Top-level domain	localdomain.local		
		vCenter Server hostname	vcenter		
		VxRail Appliance hostname	vxrail		
	Networking	ESXi starting address for IP pool	192.168.10.1		
		ESXi ending address for IP pool	192.168.10.4		
		vCenter Server IP address	192.168.10.101		
		VxRail Appliance IP address	192.168.10.100		
		Subnet mask	255.255.255.0		
		Gateway	192.168.10.254		
	Passwords	ESXi "root"			
		vCenter Server & VxRail Manager "administrator@vsphere.local"			
	Platform Controller Services (optional)	Platform controller services hostname	psscserver		
	vSphere vMotion		Starting address for IP pool	192.168.20.1	
			Ending address for IP pool	192.168.20.4	
			Subnet mask	255.255.255.0	
			VLAN ID	20	
	v SAN		Starting address for IP pool	192.168.30.1	

Table 2 Network Configuration Table (continued)

Element	Category	Description	Example	Customer value
		Ending address for IP pool	192.168.30.4	
		Subnet mask	255.255.255.0	
		VLAN ID	30	
VM Networks		VM Network name and VLAN ID	Sales / 110	
		VM Network name and VLAN ID	Marketing / 120	
		...		
		Unlimited number		
Solutions	Logging	vRealize Log Insight hostname	loginsight	
		vRealize Log Insight IP address	192.168.10.102	
		Syslog server (instead of Log Insight)		

Network planning

# CHAPTER 4

## Licensing your VxRail Appliance

Your VxRail Appliance comes with VMware vSphere pre-installed. However, you must supply vSphere licenses to activate the appliance. Your vSphere licenses can be obtained in one of two ways:

- By leveraging existing available vSphere licenses (from example, from existing ELAs)
- By acquiring new licenses from one of the following channels:
  - Dell EMC direct OEM
  - Dell EMC brokerage services
  - VMware direct
  - Partner (resellers only)

Work with your sales representative or reseller to ensure that you have the proper type and quantity of licenses available for your environment.

• Licensing overview .....	22
• Locating your partner activation code (PAC) .....	22
• Logging into the activation portal .....	22
• Preparing your vSphere license for your VxRail Appliance .....	23
• Assigning a license to your VxRail Appliance .....	24

## Licensing overview

The overall procedure for licensing your VxRail Appliance is outlined here. Refer to the subsequent sections for details of each step.

Follow these steps to activate the license on your appliance.

---

### Note

Internet access is required for this process. Customers without Internet access at their site should work with their sales representative or reseller to obtain their license.

---

### Procedure

1. Locate your partner activation code (PAC).
2. Log into the activation portal and redeem your PAC to receive your license keys.  
Use your My VMware credentials or create an account if you don't have one.
3. Prepare a vSphere license key that meets your VxRail Appliance vSphere license requirements.
4. Enter your license key on the vCenter Web Client and apply it to your VxRail Appliance.

## Locating your partner activation code (PAC)

When you buy a VxRail Appliance, you will receive a unique partner activation code (PAC) that you can use to obtain your license.

### Procedure

1. Locate your PAC.
  - A copy of the PAC is emailed to the **Bill to** and **Ship to** contacts for your order.
  - A copy of the PAC is also sent to any additional email addresses that were entered for your order.
2. Keep the PAC available as you license your appliance.

## Logging into the activation portal

Log into the activation portal to redeem your PAC and receive your license keys.

### Procedure

1. Point your Web browser to the activation portal:  
`https://www.vmware.com/oem/code.do?Name=EMC-AC`
2. Log in with your My VMware credentials.

If you do not have a My VMware account, create one using these steps:

- a. Click **Register** under **New Customers**.
- b. Enter your email address, password, and other information.



- c. Click to accept the terms and conditions and click **Continue**.

An email is sent to the address you registered with.

- d. In the email, click **Activate Now** and follow the instructions to complete your registration.

3. On the **VMware Registration for Partner Activation Codes** page, enter your PAC in the **Activation Code** box.
4. Click **Continue** and follow the instructions.

You will receive an email containing the license key. The license key must next be prepared for your VxRail Appliance, and then applied using the vSphere Web Client.

## Preparing your vSphere license for your VxRail Appliance

After you redeem your partner activation code (PAC) and receive your vSphere license, you must arrange your vSphere licenses to match your VxRail Appliance.

### Before you begin

You may need to upgrade your license key to match the product version you are using (for example, vSphere 6.0 or vSphere 6.5, and so on). Refer to the VMware knowledge base article, *How to upgrade license keys in My VMware (2006974)*.

### Procedure

1. Log into your My VMware account.
2. Navigate to the **Manage License Keys** page.
3. Select the **Combine License Keys** option.
4. Find and select the proper number of vSphere licenses to support your VxRail Appliance.

The following table lists the license support requirements. Work with your EMC or partner representative to ensure that you have the proper type and quantity of licenses available for your environment.

VxRail Appliance model	vSphere license support requirements
VxRail 60	4 hosts and 4 CPUs
VxRail 120, VxRail 120F, VxRail 160, VxRail 160F, VxRail 200, VxRail 200F, VxRail 240F, VxRail 280F	3 hosts and 6 CPUs, or 4 hosts and 8 CPUs
VxRail E460, VxRail E460F, VxRail E560, VxRail E560F, VxRail G410, VxRail G410F, VxRail P470, VxRail P470F, VxRail P570, VxRail P570F, VxRail S470, VxRail S570, VxRail V470, VxRail V470F, VxRail V570, VxRail V570F	1 to 4 hosts and 2 to 8 CPUs

5. Click **Combine**. (If your vSphere licenses support more than eight CPUs, click **Divide License Keys** instead.)
6. Follow the My VMware on-screen instructions to finish grouping the licenses.

You have the option to email your license key to yourself.

7. Create a folder called *VxRail* within My VMware with the required vSphere Enterprise Plus CPU license.

#### After you finish

For additional information on how to manage existing vSphere license keys to meet your requirements, work with your sales representative or reseller. You can also refer to the following VMware knowledgebase articles:

- How to combine license keys in My VMware (2006973)
- How to divide license keys in My VMware (2006972)
- How to upgrade license keys in My VMware (2006974)
- How to downgrade license keys in My VMware (2006975)
- Licensing ESXi 6.x and vCenter Server 6.x (2107538)

## Assigning a license to your VxRail Appliance

After combining or splitting your vSphere license and making it ready for VxRail Appliance, you must apply the license using the vCenter Web Client.

#### Procedure

1. Log into the vSphere Web Client.  
(You must have Global.Licenses privilege.)
2. Click **Home**.
3. In the **Administration** section, click **Licensing** in the left pane and select the **Licensing** tab.
4. Click **Add New Licenses** (the small plus icon: +.)
5. Enter your product license key, one license per line, and click **Next**.
6. (optional) Add a label to the license. Click **Next**.
7. Review the details and click **Finish** to add the license.
8. To assign the new license to a host, click the **Assets** tab and then click **Hosts**.
9. Select the host and click **Assign License**.  
(Shift+click to select multiple hosts to license.)  
A pop-up window shows all available licenses.
10. Select the appropriate license from the list and click **OK** to complete the licensing process.

# CHAPTER 5

## Working with VxRail Manager

Use VxRail Manager to perform administrative tasks on your VxRail Appliance cluster.

- Viewing the VxRail Manager dashboard..... 26
- Support resources.....27
- Viewing system events.....29
- Monitoring logical system health.....30
- Viewing physical system health..... 32
- Viewing VxRail Appliance system configuration and applications..... 41

## Viewing the VxRail Manager dashboard

Click **DASHBOARD** to navigate to the VxRail Manager **Dashboard** tab.

The VxRail Manager dashboard shows system health and support resources at a glance, including upgrade availability, expansion status, overall system health, support, community activity, and event history.

Figure 1 VxRail Manager Dashboard screen



### Upgrade

Upgrade displays the availability of VxRail Manager software upgrades.

### Expansion

**Expansion** shows the status additional nodes being added to your VxRail Appliance or appliance cluster.

### Overall System Health

**Overall System Health** shows the high-level system status of your VxRail Appliance. Status is shown as one of the following:

- **Healthy:** System normal. There are no major problems to address
- **Error:** An error has occurred. There is an issue that should be addressed when possible.
- **Warning:** System needs attention. There are some issues that require attention such as a disk space limit has been reached or an online support heartbeat cannot be sent.
- **Critical:** Immediate action required. There are events that must be addressed immediately to prevent downtime or data loss.

### VxRail Community

**VxRail community** shows the most recent articles and other content from the online VxRail community.

## Support

Support shows status and links to support resources, including:

- **Last Heartbeat:** The last time an EMC Secure Remote Services (ESRS) heartbeat was sent (only shown if ESRS is enabled).
- **Chat with Support:** Link to start a chat session with a support representative.
- **Open a Service Request:** Link to open a new service request ticket for support.

## Event History

Event history displays the most recent system events.

## Support resources

The VxRail Manager Support tab displays support status information, as well as support resources and links.

Figure 2 VxRail Manager Support screen



- **Last Heartbeat:** Displays the date and time of the last successful heartbeat sent by your VxRail Appliance using ESRS Connect Home. (ESRS must be running and enabled.)
- **Chat with Support:** Opens a chat session with a support representative.
- **Open a Service Request:** Navigates to eServices where you can open a service request.
- **View the Last Configuration Sent:** Displays the most recent VxRail Appliance configuration data sent using ESRS Connect Home. (ESRS must be installed and enabled.)

- **Download:** Displays available applications and add-ons available for your VxRail Appliance.
- **VxRail Community:** Displays the most recent activity from the VxRail Series community forums.
- **Knowledge Base:** Lets you search the support knowledge base and access VxRail Series support articles.

## Viewing Secure Remote Support (ESRS) information

You can verify your VxRail Appliance ESRS "heartbeat" (the last time your system communicated with the remote support service). You can also review the configuration data that was sent to ESRS.

### Before you begin

Your VxRail Appliance can use ESRS either integrated on the cluster, or through an external ESRS gateway. For integrated ESRS, it must be installed and enabled on your VxRail Appliance. You must have a support account before you activate ESRS. See [Enabling EMC Secure Remote Support \(ESRS\)](#).

### Procedure

1. Click **SUPPORT** in VxRail Manager to navigate to the **Support** tab.
2. Observe the date and time of the last ESRS communication in the **Last Heartbeat** section.
3. Click **View the Last Configuration Sent** to display the configuration data that was most recently sent to ESRS.

## Using Support eServices

Access Support eServices directly from VxRail Manager.

### Before you begin

To use Support eServices, you must have created an account. To create an account visit [support.EMC.com](http://support.EMC.com).

The VxRail Manager lets you access the following Support eServices:

- **Chat with Support:** Opens a chat session with a support representative.
- **Open a Service Request:** Opens a Web form where you can open a service request.
- **VxRail Community:** Displays the most recent activity from the VxRail Series community forums. Click a title to view the discussion.

## Chat with Support

Open a live chat session with support personnel.

### Procedure

1. Click **SUPPORT** in VxRail Manager to navigate to the **Support** tab.
2. Click **Chat with Support** to start a live chat session with a support representative.

The chat session opens. VxRail Manager transmits your appliance ID to the support representative.

## Opening a service request

Contact Customer Service to request service for your VxRail Appliance.

### Procedure

1. Click **SUPPORT** in VxRail Manager to navigate to the **Support** tab.
2. If you have more than one appliance, click the **arrow** next to **Open a Service Request**.

A list of the appliances in your cluster is shown.

3. Click the appliance for which you want service.

A service request form opens in a new browser tab. The form is prepopulated with information about the appliance you selected.

4. Fill out the required information and click **Submit**.
5. If you have a single appliance, or want to submit a general service request, click **Open a Service Request**.

A Support eServices page opens in a new browser tab where you can open a request for service.

## Viewing VxRail Series community updates

Browse and read the most recent activity from the VxRail Series community forums.

### Procedure

1. Click **SUPPORT** in VxRail Manager to navigate to the **Support** tab.
2. Observe the **VxRail Community** list.

*The most recent activity from the VxRail Series community is shown.*

3. Click the title of a message or article to view that topic in a new browser tab.

## Using the support knowledge base

Search, access, and view the support knowledge base.

### Procedure

1. Click **SUPPORT** in VxRail Manager to navigate to the **Support** tab.
2. Type your search terms in the **Knowledge Base** search field and press Enter or click the magnifying glass icon.

If you have not linked your support account with VxRail Manager, you will be prompted to log in with your support credentials.

3. Click any of the returned search results to read that knowledge base article. Or click **More...** to navigate to the knowledge base website.

## Viewing system events

The VxRail Manager **Events** tab displays a list of current system events.

- **System Events:** The System Events list displays all of the VxRail Appliance events.
- **Event Details:** Displays information for the selected event in the **System Events** list.



- If there are critical events detected, the **EVENTS** icon displays the number of unread events, in red, in the navigation bar.

Figure 3 VxRail Manager Events screen



### Procedure

1. Click **EVENTS** to navigate to the VxRail Manager Events tab.
2. Sort the events list as desired by clicking on a column heading. You can sort by ID number, Severity, or Time (including date).
3. Use the arrow buttons and scroll bars to navigate through the events list.
4. Click a row to view more information about an event.
  - New critical events are shown in red.
  - When you click an event, the red highlight is removed.
  - To set all critical events as "read", click **Mark All as Read**.
5. If a physical component is listed in the **Component** column, click its **Component ID** in the Event Details to view its status on the **Health > Physical** screen.
6. To download a list of events, click **Export all events**. An **events.csv** file is created and downloaded by your browser.

## Monitoring logical system health

You can view the health of the nodes in your VxRail Appliance cluster using the VxRail Manager **Health > Logical** screen. This screen displays CPU, memory, and storage usage for your entire cluster, individual appliances, and individual nodes.

### Procedure

1. Click **HEALTH** and then **Logical** to navigate to the VxRail Manager **Health > Logical** screen.

Figure 4 VxRail Manager Health &gt; Logical screen



The color-coded status for storage IOPS, CPU usage, and memory usage indicates the following:

- Red: Over 85% used.
- Yellow: 75 to 85% used.
- Green: Less than 75% used.

2. Click **Cluster** to view information about the cluster overall.
3. Click an appliance name to view information about that appliance.
4. Scroll to view information about the main health timeline, storage use, and nodes.
5. Click the components of a node to view more information about the Capacity (HDD, SSD), Cache (SSD), ESXi disk, or NIC.

#### Results

Use the **Health > Physical** screen to view more information about hardware status and configuration.

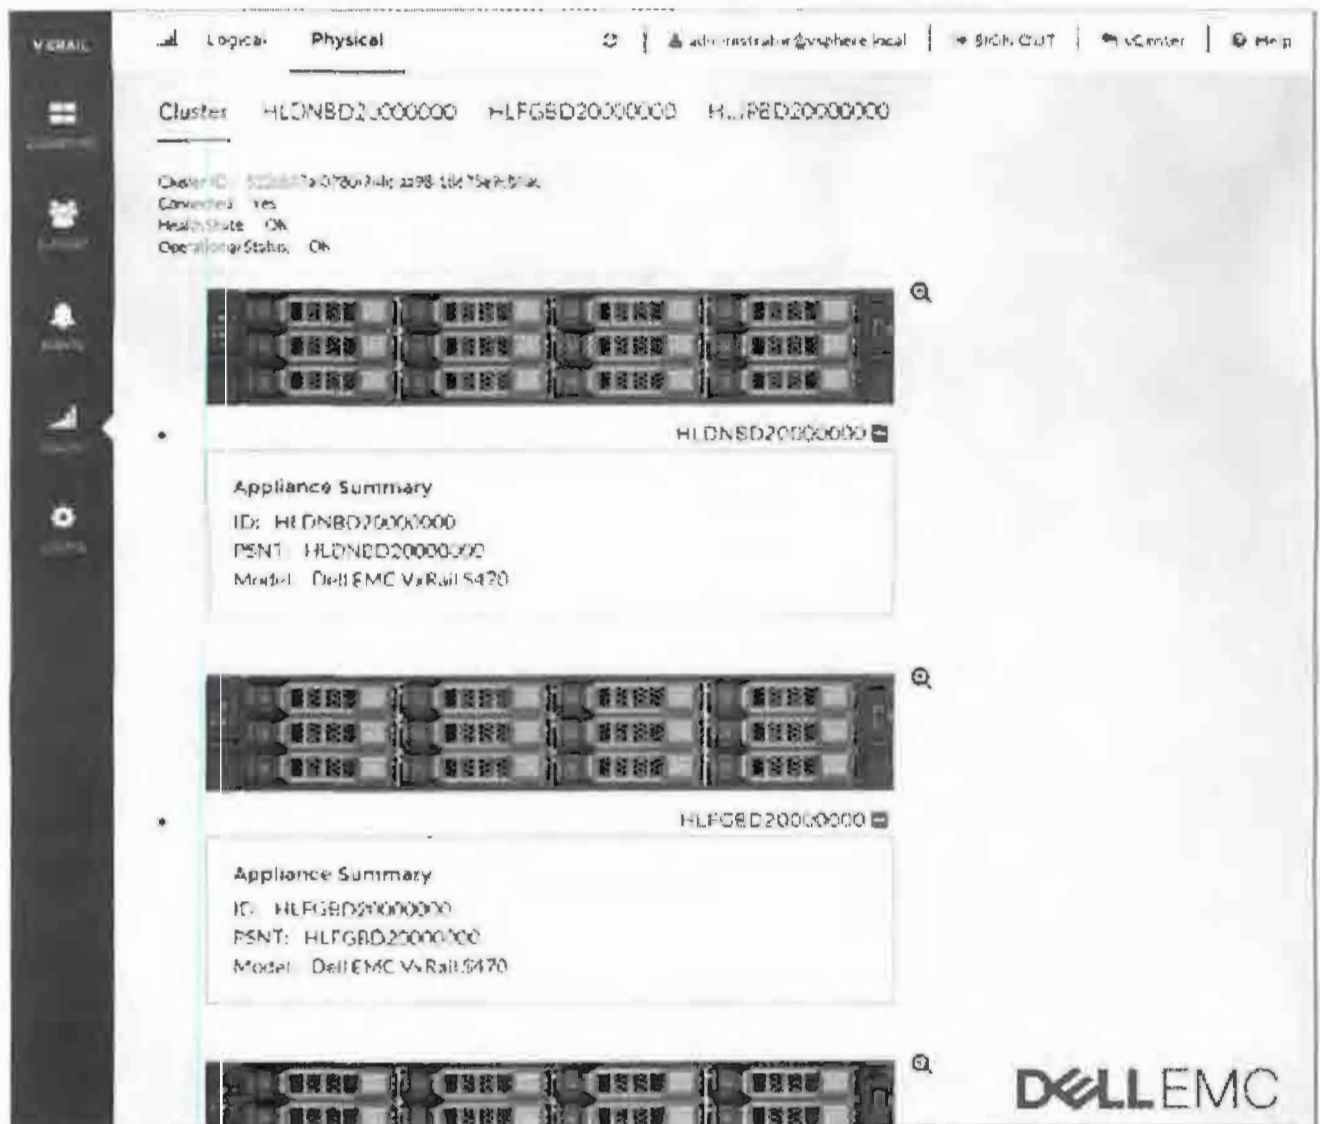
## Viewing physical system health

The **Physical** tab of the VxRail Manager **Health** window displays information about the hardware components of your appliance. A graphical representation of the appliances in your cluster makes it easy to navigate for event and status information.

Use the **Health** window **Physical** tab to view the following:

- **Appliances in the cluster:** View appliance status and information such as ID, serial number, and so on.
- **Individual appliance components:** Drill down to see status and information for appliance components such as disks, compute nodes, and power supplies.

Figure 5 VxRail Manager Health > Physical screen Cluster view



## Viewing and locating hardware events

Quickly locate and identify critical events, errors, and warnings on any appliance in the cluster.


VxRail Manager displays status icons to indicate events that occur on your appliance.

 Critical
  Error
  Warning

### Procedure

1. In VxRail Manager click **HEALTH**.  
The **Health** window is displayed.
2. Click **Physical**.  
The **Physical** tab is displayed.
3. Click **Cluster** to see a graphical overview of the appliances in the cluster.  
A picture of each appliance in the cluster is shown. Click the name below the picture to see a summary of that appliance.
4. Click an appliance name or the picture of an appliance to view more information about that appliance.  
The **Front View** and **Back View** of the appliance are displayed.

Figure 6 VxRail Manager Health > Physical screen Cluster view



5. If a status icon is displayed next to an appliance, click the appliance or the magnifying glass icon to see more information.
6. Click any appliance component to view more details.
  - Click a **Disk** in the **Front View** or **Back View** to see disk status and information.

- Click a **Node** in the **Back View** to see compute and network information.
  - Click a **power supply** in the **Back View** to see power supply status and information.
  - Click the **Back View** to see compute information.
  - Click a **NIC** in the **Back View** (E, S, P, and V models) to network information.
7. If a status icon is displayed on any component, click it to view event details in the **Health** window.
- Use your browser's back button to return to the appliance view on the **Health > Physical** tab.


## Viewing physical disk status

View status and information for any appliance disk in a cluster.

### Procedure

1. In VxRail Manager click **HEALTH**.  
The **Health** window is displayed.
2. Click **Physical**.  
The **Physical** tab is displayed.
3. Click an appliance name, picture, or the magnifying glass icon next to it.  
The **Front View** and **Back View** of the appliance are displayed.
4. Click a **Disk** in the **Front View** or **Back View**.  
The **Disk Information** panel is displayed for that disk.

Figure 7 VxRail Manager Health > Physical screen Disk Information view



5. If a status icon is displayed, click it to view event details in the **Health** window.

### Results

The **Disk Information** panel lists disk status and information about the drive, such as:


- Status LED on/off
- Serial number
- Protocol
- Enclosure and slot location
- Disk type
- Manufacturer
- Capacity
- Remaining rated write endurance level for SSDs

### Switching a disk locator LED on and off

Switch a disk's locator LED indicator on or off. The disk LED flashes to help locate the disk in an appliance.

#### Procedure

1. In VxRail Manager click **HEALTH**.  
The **Health** window is displayed.
2. Click **Physical**.  
The **Physical** tab is displayed.
3. Click an appliance name, picture, or the magnifying glass icon next to it.  
The **Front View** and **Back View** of the appliance are displayed.
4. Click a **Disk** in the **Front View**.  
The **Disk Information** panel is displayed for that node.
5. Click **Toggle Disk LED** (the wrench icon) next to the **Disk Information** panel title.



6. Select the action you want:
  - **Turn On Disk LED Indicator** starts flashing the disk LED.
  - **Turn Off Disk LED Indicator** stops flashing the disk LED.

### Checking an SSD's remaining rated write endurance

Display the remaining rated write endurance (RRWE) for an SSD in your VxRail Series E, S, P, or V model appliance.

#### Before you begin

The following VxRail models support SSD drive rated write endurance monitoring:

- VxRail E460, VxRail E460F

- VxRail P470, VxRail P470F
- VxRail S470
- VxRail V470, VxRail V470F
- VxRail E560, VxRail E560F
- VxRail P570, VxRail P570F
- VxRail S570
- VxRail V570, VxRail V570F

For SSD drives, the **Disk Information** panel displays a message when the remaining rated write endurance reaches certain levels:

- **Warning:** 30% rated write endurance remaining
- **Error:** 20% rated write endurance remaining
- **Critical:** 5% rated write endurance remaining

The Error and Critical levels are reported on the VxRail Manager **Events** screen.

Follow these steps to view the SSD remaining rated write endurance notification:

#### Procedure

1. In VxRail Manager click **HEALTH**.  
The **Health** window is displayed.
2. Click **Physical**.  
The **Physical** tab is displayed.
3. Click an appliance name, picture, or the magnifying glass icon next to it.  
The **Front View** and **Back View** of the appliance are displayed.
4. Click an **SSD** in the **Front View** or **Back View**.  
The **Disk Information** panel is displayed for that SSD.
5. Observe the remaining rated write endurance notification (if displayed).

## Viewing physical node status

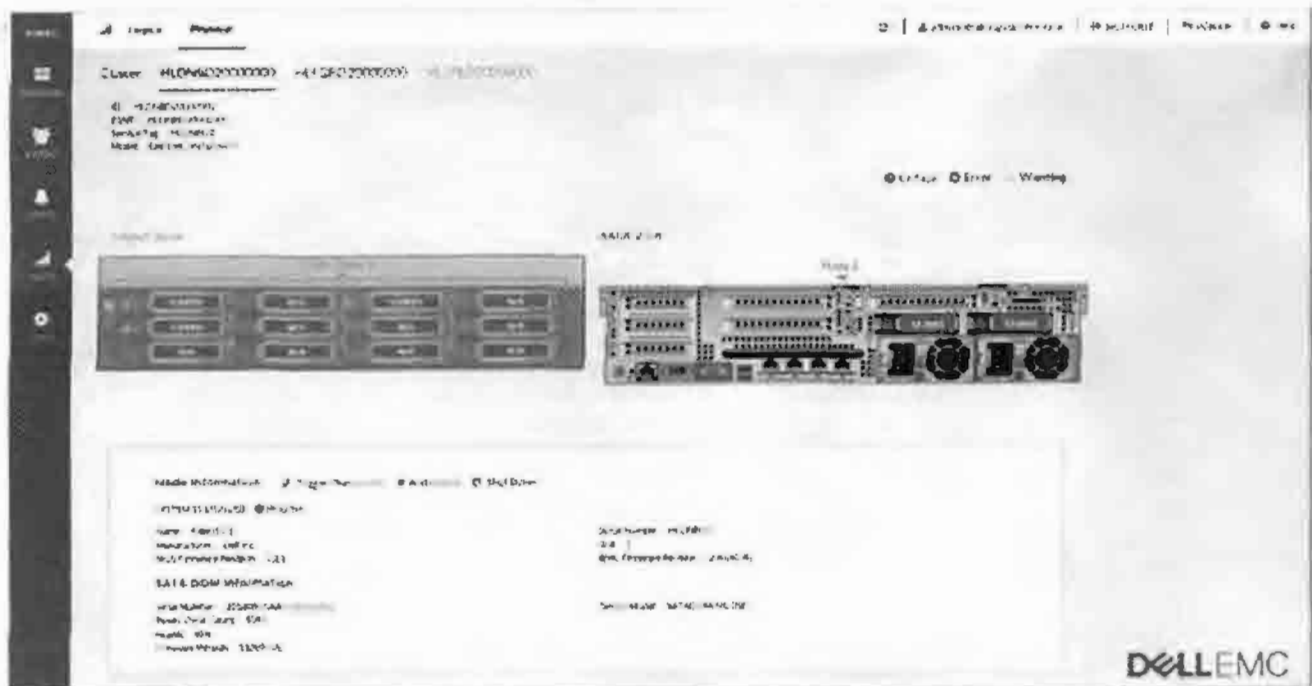
View compute, boot device, and NIC status and information for any appliance node in a cluster.

#### Procedure

1. In VxRail Manager click **HEALTH**.  
The **Health** window is displayed.
2. Click **Physical**.  
The **Physical** tab is displayed.
3. Click an appliance name, picture, or the magnifying glass icon next to it.  
The **Front View** and **Back View** of the appliance are displayed.
4. Click a chassis or a node in the **Front View** or **Back View**.  
The **Node Information** panel is displayed for that node.



Figure 8 VxRail Manager Health &gt; Physical screen Node Information view



5. If a status icon is displayed, click it to view event details in the **Events** window.

#### Results

The **Node Information** panel lists status and information about the node, such as:


- Status LED on/off
- Serial number
- BIOS firmware version
- Boot device serial number and statistics
- Network interface card (NIC) MAC addresses and status

### Switching a node chassis LED on and off

Switch a node's chassis LED indicator on or off. The chassis LED flashes to help locate the node appliance in a rack.

#### Procedure

1. In VxRail Manager click **HEALTH**.  
The **Health** window is displayed.
2. Click **Physical**.  
The **Physical** tab is displayed.
3. Click an appliance name, picture, or the magnifying glass icon next to it.  
The **Front View** and **Back View** of the appliance are displayed.
4. Click a **Node** in the **Back View**.  
The **Node Information** panel is displayed for that node.
5. Click **Toggle Chassis LED** (the wrench icon) next to the **Node Information** panel title.



6. Select the action you want:

- **Turn On Chassis LED Indicator** starts flashing the chassis LED.
- **Turn Off Chassis LED Indicator** stops flashing the chassis LED.

### Shutting down a single node

Shut down an individual node in a cluster. This is useful when replacing node hardware or performing other maintenance tasks.

The following VxRail models support shutting down a single node:

- VxRail E460, VxRail E460F
- VxRail P470, VxRail P470F
- VxRail S470
- VxRail V470, VxRail V470F
- VxRail E560, VxRail E560F
- VxRail P570, VxRail P570F
- VxRail S570
- VxRail V570, VxRail V570F

#### Procedure

1. In VxRail Manager click **HEALTH**.  
The **Health** window is displayed.
2. Click **Physical**.  
The **Physical** tab is displayed.
3. Click an appliance name, picture, or the magnifying glass icon next to it.  
The **Front View** and **Back View** of the appliance are displayed.
4. Click a **Node** in the **Back View**.  
The **Node Information** panel is displayed for that node.
5. Click **Shut Down** in the **Node Information** panel.  
The **Shut Down Node** window is displayed as VxRail Manager begins the shut down pre-check routine.
6. Correct any issues detected by the pre-check routine.
7. When the pre-check routine completes successfully, click **Continue** to shut down the node.

#### Results

During the pre-check routine, VxRail Manager checks the following:

- Checks and ensures the hosts aren't in maintenance mode
- Checks and ensures the hosts are connected

- Checks and ensures that DRS is enabled, or that all VMs on this host have been shut down
- Checks and ensures all VMs with vGPU enabled have been shut down


## Viewing power supply status

View status and information for the power supplies of any appliance disk in a cluster.

### Procedure

1. In VxRail Manager click **HEALTH**.  
The **Health** window is displayed.
2. Click **Physical**.  
The **Physical** tab is displayed.
3. Click an appliance name, picture, or the magnifying glass icon next to it.  
The **Front View** and **Back View** of the appliance are displayed.
4. Click a **power supply** in the center area of the **Back View**.  
The **Power Supply Information** panel is displayed for that power supply.

Figure 9 VxRail Manager Health > Physical screen Power Supply Information view



5. If a status icon is displayed, click it to view event details in the **Events** window.

### Results

The **Power Supply Information** panel lists status and information about the power supply, such as:

- Serial number
- Manufacturer

## Viewing physical NIC status

View NIC status and network information for your VxRail Series E, S, P, or V model appliance.

### Before you begin


The following VxRail models support NIC information monitoring in a separate NIC information panel. For the other VxRail models, NIC information is integrated in Node information panel.

- VxRail E460, VxRail E460F
- VxRail P470, VxRail P470F
- VxRail S470
- VxRail V470, VxRail V470F
- VxRail E560, VxRail E560F
- VxRail P570, VxRail P570F
- VxRail S570
- VxRail V570, VxRail V570F

### Procedure

1. In VxRail Manager click **HEALTH**.  
The **Health** window is displayed.
2. Click **Physical**.  
The **Physical** tab is displayed.
3. Click an appliance name, picture, or the magnifying glass icon next to it.  
The **Front View** and **Back View** of the appliance are displayed.
4. Click a **NIC** in the **Back View**.  
The **NIC Information** panel is displayed for that NIC.

Figure 10 VxRail Manager Health &gt; Physical screen NIC Information view



### Results

The **NIC Information** panel lists status and information about the node, such as:

- MAC address
- Link speed
- Link status
- Port number

## Viewing VxRail Appliance system configuration and applications

The VxRail Manager **Config** tab displays information and status about the physical configuration of your VxRail Appliance. **Config** also displays details about the applications installed on or available for your appliance.

- **General:** Information and settings for VxRail Manager, including:
  - VxRail Manager version
  - Linked Online Support account
  - Log collection
  - Internet proxy settings
  - Secure Remote Support (ESRS)
  - Network offline mode
  - Health monitoring suppression mode
  - System diagnostics
  - Cluster shutdown

- User interface language
- **System:** Descriptions, version information, and update controls for the system software installed on the cluster.
- **Market:** Access to qualified applications to install and run on your VxRail Appliance cluster.

## Config General tab

The **General** tab of the **Config** window displays information and settings for your VxRail Appliance.

Use the **General** tab to view the following:

- **About VxRail:** Displays the current version of VxRail Manager software running on your appliance.
- **Set Support Account:** Displays the linked Support account and allows you to link or change to a new account.
- **Log Collection:** Displays the most recent collected logs and allows you to generate a new log bundle.
- **Configure Proxy Settings:** Displays the proxy status and allows you to change, disable, or enable proxy settings for Internet connections.
- **Enable EMC Secure Remote Support (ESRS):** Displays whether or not ESRS is enabled on your VxRail Appliance.
- **Configure Network Environment Settings:** Disable certain features in offline mode.
- **Cluster Health Monitoring:** Disable system health monitoring for maintenance purposes.
- **System Diagnostics:** Run a complete system diagnostic report.
- **Shut Down Cluster:** Gracefully shut down the cluster.
- **Choose Your Language:** Select the language used in the VxRail Manager UI.

Figure 11 VxRail Manager Config &gt; General screen

The screenshot displays the VxRail Manager configuration interface, specifically the General tab. The top navigation bar includes 'System', 'Monitor', and 'General' tabs, along with user information 'Administrator@jumphere.local', 'LOG OUT', 'vCenter', and 'Help'.

**About VxRail Manager**  
VxRail Manager Version 4.0.0-4631184

**Set Support Account**  
Support User name/Password has not been set.  
**Set Support Account**

**Log Collection**  
**Generate New Log Bundle**

**Enable EMC Secure Remote Support (ESRS)**  
EMC Secure Remote Support allows your appliance to periodically send health data to support. It also allows support to connect into your appliance to correct problems, often before they become serious.  
Your current ESRS status is: Disabled  
**Enable ESRS**

**Configure Network Environment Settings**  
In offline mode, features which require Internet access will be disabled.  
Offline Mode:  On  OFF **Apply**

**Suppress Cluster Monitor**  
If suppress is on, the cluster health status will not be monitored. This is can be useful when part of the cluster is disabled, such as for node replacement.  
Suppress Mode:  On  OFF **Apply**

**System Diagnostics**  
Diagnostics for the VxRail cluster health status.  
**Diagnostics**

**Shut Down Cluster**  
Shut down the entire cluster.  
**Shut Down**

**Choose Your Language**

ENGLISH	FRANÇAIS	DEUTSCH
한국어	日本語	简体中文

**DELL EMC**

## Viewing the VxRail Manager version number

Display the current version of VxRail Manager software running on your appliance.

### Procedure

1. In VxRail Manager click **CONFIG**.  
The **Config** window is displayed.
2. Click **General**.  
The **General** tab is displayed.
3. Observe **About VxRail Manager** for the VxRail Manager software version.

## Linking your Online Support account to VxRail Manager

For convenience, you can link your Online Support account with VxRail Manager and access support resources without having to log in separately.

### Before you begin

You must have an account before you can link it with VxRail Manager. To create a support account, see *Registering for online support*.

### Procedure

1. In VxRail Manager click **CONFIG**.  
The **Config** window is displayed.
2. Click **General**.  
The **General** tab is displayed.
3. Click one of the following buttons:
  - **Set Support Account** (if no account has been set)
  - **Change Support Account** (if an account was previously set)The **Username** and **Password** fields are displayed.
4. Type your account credentials in the **Username** and **Password** fields.  
Leave both fields blank if you wish to unlink all accounts.
5. Click **Login**.  
VxRail Manager will log into your account and display that your account is in use.

## Generating a diagnostic log bundle

Create and download a diagnostics log bundle for your cluster.

VxRail Manager compiles a diagnostic log of its own operations. It does not include log information for the entire cluster.

### Procedure

1. In VxRail Manager click **CONFIG**.  
The **Config** window is displayed.
2. Click **General**.  
The **General** tab is displayed.
3. In the **Log Collection** section click **Generate New Log Bundle**.  
VxRail Manager will gather diagnostics information and create a TAR file.



4. Click **Open** to view the file or **Save** to store it on your local machine.

### Configuring a network proxy server

Configure proxy settings to allow connections to an external network through your proxy server. You can update, enable, and disable the proxy settings.

#### Before you begin

You must supply an existing proxy server. The type of the proxy should be either HTTP or SOCK4/5.

The following table lists the supported proxy types for VxRail Manager and ESRS VE.

**Table 3** Supported proxy types

Proxy Type	VxRail Manager	ESRS VE
SOCKS4 without authentication	Supported	Not supported
SOCKS5 without authentication	Supported	Supported
SOCKS5 with authentication	Supported	Not supported
HTTP without authentication (anonymous)	Supported	Supported
HTTP with authentication (basic authentication)	Supported	Supported

The proxy is used for the Internet connection to the EMC community, knowledge base search, appliance market download, and so on.

The proxy configuration is applied to the ESRS virtual environment during ESRS enablement and activation. After ESRS is activated, the proxy information can be set from the ESRS side independently but it will always be overwritten if new a proxy is set from VxRail Manager.

#### Procedure

1. In VxRail Manager click **CONFIG**.  
The **Config** window is displayed.
2. Click **General**.  
The **General** tab is displayed.
3. Verify the value of **Your current proxy status**:
  - **Disabled**: Proxy is not enabled.
  - **Enabled**: Proxy is enabled.
4. Do one of the following:
  - If your proxy status is **Disabled**, and you need to configure the proxy, proceed to step 5.
  - If your proxy status is **Enabled**, and you need to update the proxy, skip to step 8.
  - If your proxy status is **Enabled**, and you need to disable the proxy, skip to step 9.
5. Click **Enable Proxy**.  
The proxy configuration form expands.

6. Enter the information for the proxy you want to configure:
  - a. Select the expected proxy type from the **Type** dropdown box.  
The supported proxy types includes **HTTP** and **SOCKS**.
  - b. If **SOCKS** is selected, choose **SOCKS4** or **SOCKS5** as the **SOCKS** type.
  - c. In **Address** field, enter the **IP/Hostname** and **Port** number of the proxy server you are using.
  - d. If you choose **HTTP** or **SOCKS5** in step a, you can provide credentials for accessing the proxy server by checking **Authenticate using following credentials**. Provide the username and password in the corresponding fields.
  - e. If **ESRS** is enabled, enter the **ESRS** admin password in the **ESRS password** field.

7. Click **Enable**.

The configured proxy settings are enabled for VxRail Manager.

When your proxy status is **Enabled**, the process is complete. A message "Proxy settings have been enabled" is displayed.

8. To modify the proxy settings, do the following:
  - a. Click **Update Proxy**.  
The proxy configuration form is expanded.
  - b. Modify the fields you want to update in the form.
  - c. If **ESRS** is enabled, input the **ESRS** admin password in the **ESRS password** field.
  - d. Click **Update**.

The configured proxy settings are changed and applied for VxRail Manager.

9. To disable the proxy, do the following:
  - a. Click the triangle button beside the **Update Proxy** button.
  - b. Click **Disable Proxy** in the dropdown box.
  - c. If **ESRS** is enabled, input the **ESRS** admin password in the **ESRS password** field and click **Disable Proxy**.

When your proxy status is **Disabled**, the process is complete. A message "Proxy settings have been disabled" is displayed.

## Enabling Secure Remote Support (ESRS)

ESRS provides secure, automated access between Support and your VxRail Appliance. You must install and activate an integrated ESRS or register to an External ESRS gateway to enable some features in VxRail Manager.

You must have a support account before you enable ESRS. To create a support account, see *Registering for online support*.

---

### Note

If you plan to set up ESRS, your Online Support account must be linked to VxRail Manager under the same party ID or it will not work properly.

---

## Enabling integrated Secure Remote Support (ESRS)

Follow these steps to activate integrated ESRS on your VxRail Appliance.

### Procedure

1. In VxRail Manager click **CONFIG**.  
The **Config** window is displayed.
2. Click **General**.  
The **General** tab is displayed.
3. Verify the value of **Your current ESRS status is**.
  - **Disabled**: ESRS is not activated.
  - **Deployed**: ESRS has been installed but not activated.
  - **Enabled**: ESRS is activated.
4. Do one of the following:
  - If your ESRS status is **Disabled**, proceed to step 5.
  - If your ESRS status is **Deployed**, skip to step 11.
  - If your ESRS status is **Enabled**, skip to step 14.
5. Click **Enable ESRS**.  
The **Enable ESRS** window is displayed.
6. Uncheck **Join External ESRS**.  
The fields for configuring integrated ESRS are displayed.
7. In the **IP** field, enter a valid IP address on the management network used by your VxRail Appliance.  
  
ESRS must be on the same management network as vCenter and VxRail Manager to receive alerts and broadcast them back to Support.
8. Verify that the site ID and contact information are correct.  
  
These fields are pre-filled based on your Support account. However, you can manually update them if you wish.
9. Set the passwords for the root account and the admin account.
10. Click **Submit**.  
  
A status bar is displayed as ESRS is deployed on your system. When the process is complete, an access code is emailed to the address specified for your support account.
11. When your ESRS status is **Deployed**, click **Activate ESRS**.  
The **Enable ESRS** window is displayed.
12. Type or paste the access code in the **Access code** field.  
  
The access code is valid for 30 minutes. If your code has expired, click **Regenerate access code** to receive a new code by email.
13. Click **Submit**.  
  
A status bar is displayed as ESRS is activated on your system. When the process is complete, your ESRS status shows **Enabled**.
14. When your ESRS status is **Enabled**, the process is complete.

15. If you want to reset ESRS on your system, click **Reset ESRS**. In the **Reset ESRS** dialog, click **Reset** to confirm.

You can re-enable ESRS by following the steps in this procedure.

### Enabling external Secure Remote Support (ESRS)

Follow these steps to configure external ESRS on your VxRail Appliance.

#### Procedure

1. In VxRail Manager click **CONFIG**.  
The **Config** window is displayed.
2. Click **General**.  
The **General** tab is displayed.
3. Verify the value of **Your current ESRS status is**.
  - **Disabled**: ESRS is not activated.
  - **Deployed**: ESRS has been installed but not activated.
  - **Enabled**: ESRS is activated.
4. Do one of the following:
  - If your ESRS status is **Disabled**, proceed to step 5.
  - If your ESRS status is **Deployed**, skip to step 11.
  - If your ESRS status is **Enabled**, skip to step 11 or 12.
5. Click **Enable ESRS**.  
The **Enable ESRS** window is displayed.
6. Check **Join External ESRS**.  
The fields for configuring external ESRS are displayed.
7. In the **IP** field, enter the external ESRS gateway IP address for VxRail Manager registration.
8. Verify that the site ID and contact information are correct.
9. Add another external ESRS for failover.
10. Click **Submit**.  
VxRail Manager is registered with the specified external ESRS.  
When your ESRS status shows as **Enabled**, the process is complete.
11. To view or reconfigure external ESRS on your system, follow these steps:
  - a. Click **Show/Reconfig ESRS**.  
The current registered ESRS is displayed.
  - b. Make your changes.
  - c. Click **Submit**.  
The ESRS configuration is modified.
12. To reset ESRS on your system, follow these steps:
  - a. Click **Reset ESRS**.
  - b. In the **Reset ESRS** dialog, click **Reset** to confirm.

You can re-enable ESRS by following the steps in this procedure.

## Setting Internet online/offline mode (Network Environment Settings)

Follow these steps to enable or disable your VxRail Appliance's connection to the Internet.

---

### Note

Some VxRail Manager features are disabled when the appliance Internet connection is disabled.

---

### Procedure

1. In VxRail Manager click **CONFIG**.  
The **Config** window is displayed.
2. Click **General**.  
The **General** tab is displayed.
3. In the **Configure Network Environment Settings** section, select one of the following:
  - **Internet Connection Connected:** Internet access for your appliance is enabled.
  - **Internet Connection Disconnected:** Internet access for your appliance is disabled.
4. Click **Apply**.

### Results

When **Internet Connection Disconnected** is selected, some VxRail Manager features are disabled. A message is displayed for the disabled features.

## Enabling and disabling cluster health monitoring

Follow these steps to set whether or not VxRail Manager monitors the system health of the cluster.

Cluster health monitoring is normally enabled. You may want to disable it when you are performing service, such as replacing disks and so on.

### Procedure

1. In VxRail Manager click **CONFIG**.  
The **Config** window is displayed.
2. Click **General**.  
The **General** tab is displayed.
3. In the **Suppress Cluster Monitor** section, select one of the following:
  - **Health Monitoring On:** Cluster health is monitored and reported. This is the default setting.
  - **Health Monitoring Muted:** Cluster health is not monitored.
4. Click **Apply**.

### Results

When cluster health monitoring is disabled (**Health Monitoring** set to **Muted**) VxRail Manager displays a banner indicating that monitoring is muted. You can close the banner by clicking the **X** on the right side of the banner.

## Running a complete system diagnostic

Follow these steps to run a complete system diagnostic report for your VxRail Appliance.

### Procedure

1. In VxRail Manager click **CONFIG**.  
The **Config** window is displayed.
2. Click **General**.  
The **General** tab is displayed.
3. In the **System Diagnostics** section, click **Diagnostic**.  
VxRail Manager runs a system-wide diagnostic analysis of the appliance. The **System Diagnostics** window is displayed.
4. Observe the **Health check list** in the **System Diagnostics** window.  
Click items in the list to expand them for more information.
5. When you are finished with the diagnostic report click **Close**.

## Shutting down a VxRail Appliance cluster

You can shut down your VxRail Appliance cluster from VxRail Manager. This feature provides a graceful, automated process for properly shutting down a cluster.

When you shut down a cluster, VxRail Manager automatically does the following:

- Shuts down related virtual machines and services
- Performs system health diagnostics and maintenance mode diagnostics
- Indicates any errors or conditions that prevent shutting down

Follow these steps to shut down a VxRail Appliance cluster:

### Procedure

1. In VxRail Manager click **CONFIG**.  
The **Config** window is displayed.
2. Click **General**.  
The **General** tab is displayed.
3. In the **Shut Down Cluster** section, click **Shut Down**.  
The **Shut Down Cluster** dialog box is displayed.
4. Click **Confirm** to continue shutting down.  
VxRail Manager begins a precheck of the system and displays the progress on the **Shut Down Pre-Check list**.
5. If the shutdown precheck process fails, VxRail Manager lists error messages in the **Pre-check Result** column. Resolve any issues and click **Retry**.  
You can also check **ignore and continue** and then click **Shut Down** by force.

6. When the shutdown precheck process successfully completes, click **Shutdown**.  
A progress bar is displayed as the automated shut down process continues.
7. Click **Close** to exit.  
Wait for all the LEDs to turn off.

#### Results

When the shut down process is complete, VxRail Manager indicates that shut down is successful.

If the process does not complete successfully, follow the on-screen prompts to correct any errors.

### Selecting a language for VxRail Manager

You can specify the language that displays in the VxRail Manager user interface.

The VxRail Manager user interface supports the following languages:

- English (default)
- French
- German
- Korean
- Japanese
- Simplified Chinese

#### Procedure

1. In VxRail Manager click **CONFIG**.  
The **Configuration** window is displayed.
2. Click **General**.  
The **General** tab is displayed.
3. In the **Choose Your Language** section of the screen, click the language you want to use.  
The VxRail Manager user interface will immediately reload in the language you selected.

## Config System tab

The **System** tab of the VxRail Manager **Config** window displays information about the applications and add-ons currently installed on your VxRail Appliance. It also allows you to download updates for those components.

Figure 12 VxRail Manager Config > System tab



Use the **System** tab to view the following:

- **System Software Installed on this Cluster:** The list displays the name, description, and installed components for the software and add-ons installed on your VxRail Appliance.
- **Software updates:** If an update is available for an application, an **Internet Upgrade** button is displayed.

When a new version of an application is available, **CONFIG** in the left navigation bar displays a highlighted number.

### Viewing installed applications

Use the **System** tab on the **Config** window to view a list of installed system software.

#### Procedure

1. In VxRail Manager click **CONFIG**.  
The **Config** window is displayed.
2. Click **System**.  
The **System** tab is displayed.
3. Scroll the **System Software Installed on this Cluster** list, if necessary, to view system software installed on your VxRail Appliance.  
The list includes a description and version number for each application installed. If a newer version of an application is available, the **Internet Upgrade** button is also displayed.
4. To view the components of the installed system software, click **Installed Components**.





The VxRail Manager, vCenter, ESXi, VxRail VIB, and PTAgent (for Dell PowerEdge appliances) versions are displayed.

#### Results

- If the installed software version is different from the factory version or the last upgraded version, VxRail Manager shows "The version is <xxx> (supported, reference version is <yyy>)". Related system events are generated.
- If the installed software version is not supported, VxRail Manager shows "The version is <xxx> (not supported). See KB500665". Related system events are generated.
- If the installed software version is incompatible with other software, VxRail Manager shows "The version is <xxx> (incompatible with <yyy>). See KB500666". Related system events are generated.

## Updating system software

Use the **System** tab on the **Config** window to install updates for the system software installed on your VxRail Appliance.

When a new version is available, **CONFIG** in the left navigation bar displays a highlighted number. The **Internet Upgrade** button is also displayed for the application in the **System Software Installed on this Cluster** list.

---

#### Note

You may not be able to update your system software in some circumstances:

- If you are using a stretch cluster VxRail configuration
- If you have standard vSphere licenses
- If your cluster is in an unhealthy state or has critical health alarms

Open a service request, or contact your sales representative or reseller to arrange help in updating your system.

---

#### Procedure

1. In VxRail Manager click **CONFIG**.  
The **Config** window is displayed.
2. Click **System**.  
The **System** tab is displayed.
3. Scroll the **System Software Installed on this Cluster** list, if necessary, to view the system software installed on your VxRail Appliance.

If a newer version is available, the **Internet Upgrade** button is displayed.

4. Click the upgrade button to view the upgrade options:

- **Local Upgrade:** Allows you to upload a new version of the software from your local storage. This is useful if you do not have Internet access from your VxRail Appliance.
- **Internet Upgrade:** Allows you to download a new version of the software. You must have Internet access to download system software.

Not all applications will show both options.

5. Click the option you want.

The software is uploaded to the VSAN.

6. Click **Continue** to begin upgrading the software.

7. Enter your login information for VxRail Manager and for vCenter Server, the PSC root account and password, and the vCenter Server administrator account and password and then click **Submit**.

VxRail Manager performs a series of system checks and begins the update process.

Enter information for VxRail Manager and vCenter Server , PSC root account and password and vCenter Server administrator account and password and then click **Submit**

The upgrade process status is displayed. When the process is done, VxRail Manager displays "Upgrade complete". If any problems are encountered during installation, VxRail Manager will display an appropriate message and steps to correct the issue.

8. Click **Refresh** to reload the **System Software Installed on this Cluster** screen.

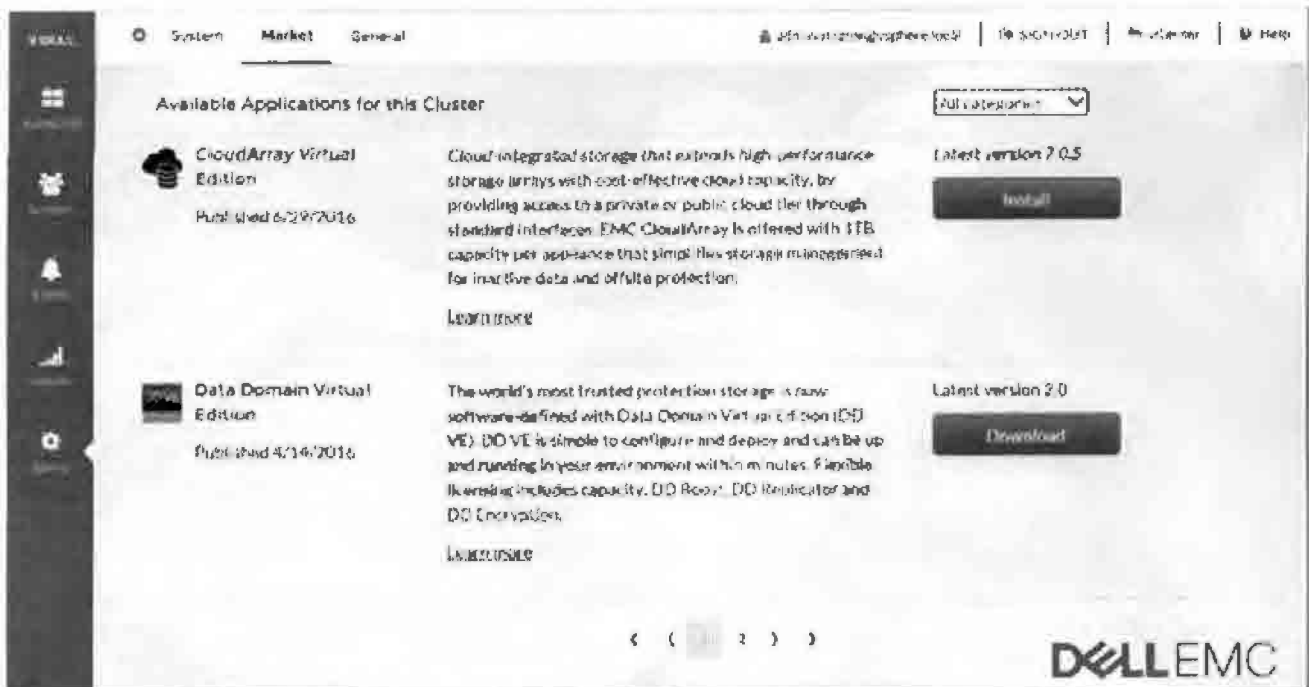
When the installation is complete, the new version is displayed in the **Installed Components** list.

## Using the VxRail Market

The VxRail Market lets you download, install, and upgrade qualified software products for your appliance.

Choose from a list of applications that add functionality, protection, or management to your VxRail Appliance.

Figure 13 VxRail Market Config &gt; Market tab



### Procedure

1. In VxRail Manager click **CONFIG**.  
The **Config** window is displayed.
2. Click **Market**.  
The **Market** tab is displayed.
3. Scroll the **Available Applications for this Cluster** list, if necessary, to view all the applications available for your VxRail Appliance.  
The list includes a description and version number for each application.
4. Filter the application list using the filter selector at the top of the list, if desired.
5. Click **Learn more** if you want to view information about the application.  
The application page opens in a separate browser tab.
6. In VxRail Manager, do one of the following to install the application on your appliance:
  - Click **Install** to install an application directly.
  - Click **Download** to navigate to an external Web page where you can download and install the application.

Refer to Installing an application from the VxRail Market. You can install multiple instances of an application.
7. Multiple instances of an application may be installed. To view and manage instances do the following:
  - a. Click the arrow next to the name of the application to expand the instance listing.

If there is only one instance of an application, the expand arrow is not displayed.

- b. Scroll to the instance you want to manage.
- c. View and manage the instance of the application as desired.

The following status messages may be displayed for an application instance:

- **Downloading:** The application is currently downloading.
- **Pending:** The application is waiting until another application finishes downloading.
- **IP address:** The application is installed.

## Installing an application from the VxRail Market

Install a qualified software product from the VxRail Market.

### Procedure

1. In VxRail Manager click **CONFIG**.  
The **Config** window is displayed.
2. Click **Market**.  
The **Market** tab is displayed.
3. Scroll the **Available Applications for this Cluster** list to find the application you want.
4. Click **Install** to install the application.
5. Fill out the fields as appropriate for the application being installed. For example, you may enter:

- **Virtual Machine Name**
- **IP Management address**
- Time zone

The fields vary depending on the application. Some fields may be pre-populated.

6. Click **Install**.  
A status bar shows the progress of the installation. The status messages include:
  - **Downloading:** The application is currently downloading.
  - **Pending:** The application is waiting until another application finishes downloading.
  - **IP address:** The application is installed.
7. To install multiple instances of an application, repeat this procedure.

## Managing applications and add-ons with VxRail Manager

Manage your installed software products from the VxRail Market.

### Procedure

1. In VxRail Manager click **CONFIG**.  
The **Config** window is displayed.
2. Click **Market**.

The **Market** tab is displayed.

3. Scroll the **Available Applications for this Cluster** list to find the application you want.

Filter the application list using the filter at the top of the list, if desired.

4. If there are multiple instances of the application, follow these steps:
  - a. Click the arrow next to the name of the application to expand the instance listing.

If there is only one instance of an application, the expand arrow is not displayed.

- b. Scroll to the instance you want to manage.
5. For the selected application, do either of the following:
  - Click **Open** to view the VM console.
  - Click **Manage** to view the application's management interface. (Not every application has a management option.)



# CHAPTER 6

## Other administrative tasks

You can perform various administrative tasks outside of VxRail Manager. Not every task is supported on every version of the VxRail Appliance. Instruction for how to perform these tasks are found in either the Solve Desktop application or in the VMware vSphere documentation.

The following sections list the administrative tasks you can perform, and where to find the instructions.

- Active Directory user management tasks..... 60
- Cluster management tasks..... 60
- Data center management tasks..... 60
- Datastore management tasks..... 61
- DNS management tasks..... 61
- ESRS management tasks..... 62
- ESXi host management tasks..... 62
- External vCSA and PSC management tasks..... 62
- Internal vCSA and PSC management tasks..... 63
- Migrating from a VxRail vCenter Server to a customer-supplied vCenter Server  
..... 64
- iDRAC configuration management tasks..... 65
- iSCSI management tasks..... 65
- Log server management tasks..... 66
- NFS management tasks..... 66
- NTP server management tasks..... 66
- vSphere Distributed Switch (VDS) management tasks..... 66
- VM folder management tasks..... 68
- vSphere Standard Switch (VSS) management tasks..... 68
- VxRail Manager application management tasks..... 69

## Active Directory user management tasks

For step-by-step instructions on managing AD users, refer to the SolVe Desktop application.

The following table lists the AD user administrative tasks by VxRail Manager software release.

Table 4 AD user management tasks

Task	4.0.1xx	4.0.200	4.0.3xx	4.0.400	4.5.x	Notes
AD Authentication in External vCenter Server	Yes	Yes	Yes	Yes	Yes	Refer to the SolVe Desktop application
Join an Active Directory Domain	Yes	Yes	Yes	Yes	Yes	Refer to the SolVe Desktop application

## Cluster management tasks

For step-by-step instructions on managing clusters, refer to the SolVe Desktop application or VMware vSphere documentation.

The following table lists the cluster administrative tasks by VxRail Manager software release.

Table 5 Cluster management tasks

Task	4.0.1xx	4.0.200	4.0.3xx	4.0.400	4.5.x	Notes
Change DRS Setting	Yes	Yes	Yes	Yes	Yes	Refer to vSphere documentation, "Disable DRS"
Change the VxRail Cluster EVC Mode	Yes	Yes	Yes	Yes	Yes	Refer to the SolVe Desktop application
Change HA Setting	Yes	Yes	Yes	Yes	Yes	Refer to vSphere Documentation, "vCenter High Availability"
Rename the VxRail Cluster	Yes	Yes	Yes	Yes	Yes	Refer to the SolVe Desktop application

## Data center management tasks

For step-by-step instructions on managing data center settings, refer to the SolVe Desktop application.

The following table lists the data center administrative tasks by VxRail Manager software release.



Table 6 Data center management tasks

Task	4.0.1xx	4.0.200	4.0.3xx	4.0.400	4.5.x	Notes
Rename the VxRail Data Center (customer-supplied vCenter)	Not supported	Yes	Yes	Yes	Yes	Refer to the Solve Desktop application
Rename the VxRail Data Center (VxRail-supplied vCenter)	Not supported	Yes	Yes	Yes	Yes	Refer to the Solve Desktop application

## Datastore management tasks

For step-by-step instructions on managing datastores, refer to the *Solve Desktop* application.

The following table lists the datastore administrative tasks by VxRail Manager software release.

Table 7 Datastore management tasks

Task	4.0.1xx	4.0.200	4.0.3xx	4.0.400	4.5.x	Notes
Rename the VxRail Datastore	Yes	Yes	Yes	Yes	Yes	Refer to the Solve Desktop application

## DNS management tasks

For step-by-step instructions on managing DNS servers, refer to the *Solve Desktop* application.

The following table lists the DNS administrative tasks by VxRail Manager software release.

Table 8 DNS management tasks

Task	4.0.1xx	4.0.200	4.0.3xx	4.0.400	4.5.x	Notes
Change VxRail Cluster's DNS IP Address (customer-supplied vCenter)	Not supported	Yes	Yes	Yes	Yes	Refer to the Solve Desktop application
Change VxRail Cluster's DNS IP Address (VxRail-supplied vCenter)	Yes	Yes	Yes	Yes	Yes	Refer to the Solve Desktop application

## ESRS management tasks

After ESRS is installed and activated on your VxRail appliance, you cannot change the parameters listed in the following table.

The following table lists the ESRS administrative tasks by VxRail Manager software release.

Table 9 ESRS management tasks

Task	4.0.1xx	4.0.200	4.0.3xx	4.0.400	4.5.x	Notes
Delete ESRS	Not supported	Not supported	Not supported	Not supported	Not supported	
Change the ESRS IP address	Not supported	Not supported	Not supported	Not supported	Not supported	
Change the ESRS hostname	Not supported	Not supported	Not supported	Not supported	Not supported	
Rename the ESRS VM	Not supported	Not supported	Not supported	Not supported	Not supported	

## ESXi host management tasks

For step-by-step instructions on managing ESXi hosts, refer to the Solve Desktop application.

The following table lists the ESXi host administrative tasks by VxRail Manager software release.

Table 10 ESXi host management tasks

Task	4.0.1xx	4.0.200	4.0.3xx	4.0.400	4.5.x	Notes
Change the ESXi host name and IP Address	Not supported	Yes	Yes	Yes	Yes	Contact Support
Change the VxRail External Management User Password	Yes	Yes	Yes	Yes	Yes	Refer to the Solve Desktop application
Change the VxRail ESXi Root Password	Yes	Yes	Yes	Yes	Yes	Refer to the Solve Desktop application

## External vCSA and PSC management tasks

For step-by-step instructions on managing the vCenter Server Virtual Appliance (vCSA) and the Platform Services Controller (PSC) on VxRail systems with an external vCenter, refer to the Solve Desktop application or VMware vSphere documentation.

The following table lists the vCSA and PSC administrative tasks for external vCenter systems by VxRail Manager software release.

Table 11 vCSA and PSC (external vCenter) management tasks

Task	4.0.1xx	4.0.200	4.0.3xx	4.0.400	4.5.x	Notes
Change the PSC IP address	Not supported	Not supported	Not supported	Not supported	Not supported	
Change the vCSA IP address	Not supported	Not supported	Not supported	Not supported	Not supported	
Rename the PSC VM	Yes	Yes	Yes	Yes	Yes	Refer to vSphere Documentation, "Changing a Virtual Machine Name"
Rename the PSC host	Not supported	Not supported	Not supported	Not supported	Not supported	
Rename the vCSA host	Not supported	Not supported	Not supported	Not supported	Not supported	
Rename the vCSA VM	Yes	Yes	Yes	Yes	Yes	Refer to vSphere Documentation, "Changing a Virtual Machine Name"
Change the password for administrator@vsphere.local	Not supported	Yes	Yes	Yes	Yes	Refer to vSphere Documentation, "Change your vCenter Server Single Sign-on Password"
Change the VxRail External Management User Password (management@localos)	Yes	Yes	Yes	Yes	Yes	Refer to the SolVe Desktop application
Change the password of the PSC root user (root@localos)	Yes	Yes	Yes	Yes	Yes	Refer to vSphere Documentation, "Change the Password of the Root User"
Change the password of the vCSA root user (root@vCSA)	Yes	Yes	Yes	Yes	Yes	Refer to vSphere Documentation, "Change the Password of the Root User"

## Internal vCSA and PSC management tasks

For step-by-step instructions on managing the vCenter Server Virtual Appliance (vCSA) and the Platform Services Controller (PSC) on VxRail systems with an internal vCenter, refer to the SolVe Desktop application or VMware vSphere documentation.

The following table lists the vCSA and PSC administrative tasks for internal vCenter systems by VxRail Manager software release.

Table 12 vCSA and PSC (internal vCenter) management tasks

Task	4.0.1xx	4.0.200	4.0.3xx	4.0.400	4.5.x	Notes
Change the VxRail vSphere Components SSL Certificates for Internal vCSA	Yes	Yes	Yes	Yes	Yes	Refer to the Solve Desktop application
Change the PSC IP address	Yes	Yes	Yes	Yes	Yes	Contact Support
Change the vCSA IP address	Yes	Yes	Yes	Yes	Yes	Contact Support
Rename the PSC VM	Not supported	Not supported	Yes	Yes	Yes	Refer to the Solve Desktop application
Rename the PSC host	Not supported	Not supported	Not supported	Not supported	Not supported	
Rename the vCSA host	Not supported	Not supported	Not supported	Not supported	Not supported	
Rename the vCSA VM	Not supported	Not supported	Yes	Yes	Yes	Refer to the Solve Desktop application
Change the Internal vCSA Administrator Password (administrator@vsphere.local)	Not supported	Yes	Yes	Yes	Yes	Refer to the Solve Desktop application
Change the VxRail External Management User Password (management@localos)	Yes	Yes	Yes	Yes	Yes	Refer to the Solve Desktop application
Change the password of the PSC root user (root@localos)	Not supported	Yes	Yes	Yes	Yes	Refer to the Solve Desktop application
Change the password of the vCSA root user (root@vCSA)	Not supported	Yes	Yes	Yes	Yes	Refer to the Solve Desktop application

## Migrating from a VxRail vCenter Server to a customer-supplied vCenter Server

For assistance with migrating a VxRail cluster from a VxRail-based vCenter Server to a customer-supplied (external) vCenter Server contact Support.

The following table lists the vCenter migration tasks by VxRail Manager software release.

Table 13 vCenter migration tasks

Task	4.0.1xx	4.0.200	4.0.3xx	4.0.400	4.5.x	Notes
Migrate VxRail Cluster from Embedded vCenter to External vCenter	Not supported	Not supported	Yes	Yes	Yes	Contact Support

## iDRAC configuration management tasks

After your VxRail Appliance is installed and run for the first time, be cautious when changing any integrated Dell Remote Access Controller (iDRAC) settings.

The following table lists the iDRAC administrative tasks by VxRail Manager software release.

Table 14 iDRAC management tasks

Task	4.0.1xx	4.0.200	4.0.3xx	4.0.400	4.5.x	Notes
Change the OS to iDRAC passthrough via the USB setting	Not supported	Not supported	Not supported	Not supported	Not supported	Do not change this setting
Reset iDRAC to factory defaults	Not supported	Not supported	Not supported	Not supported	Not supported	Do not reset iDRAC to factory defaults
Modify the vSwitch for the iDRAC network	Not supported	Not supported	Not supported	Not supported	Not supported	Do not change any settings such as vusb, vnic, vmkernel port, or iDRAC standard portgroup

## iSCSI management tasks

For step-by-step instructions on connecting to an iSCSI, refer to the External Storage section of this Administration Guide and the VMware vSphere documentation.

The following table lists the iSCSI administrative tasks by VxRail Manager software release.

Table 15 iSCSI management tasks

Task	4.0.1xx	4.0.200	4.0.3xx	4.0.400	4.5.x	Notes
Connect to an iSCSI server	Yes	Yes	Yes	Yes	Yes	External Storage section of this Administration Guide

## Log server management tasks

For step-by-step instructions on managing the log server, refer to the SolVe Desktop application.

The following table lists the log server administrative tasks by VxRail Manager software release.

Table 16 Log server management tasks

Task	4.0.1xx	4.0.200	4.0.3xx	4.0.400	4.5.x	Notes
Point to a new external log server IP address	Not supported	Yes	Yes	Yes	Yes	Refer to the SolVe Desktop application

## NFS management tasks

For step-by-step instructions on connecting to an NFS server, refer to the External Storage section of this Administration Guide and the VMware vSphere documentation.

The following table lists the NFS administrative tasks by VxRail Manager software release.

Table 17 NFS management tasks

Task	4.0.1xx	4.0.200	4.0.3xx	4.0.400	4.5.x	Notes
Connect to an NFS server	Yes	Yes	Yes	Yes	Yes	External Storage section of this Administration Guide

## NTP server management tasks

For step-by-step instructions on managing network time protocol (NTP) servers, refer to the SolVe Desktop application.

The following table lists the NTP server administrative tasks by VxRail Manager software release.

Table 18 NTP server management tasks

Task	4.0.1xx	4.0.200	4.0.3xx	4.0.400	4.5.x	Notes
Point to a new NTP server IP address	Yes	Yes	Yes	Yes	Yes	Refer to the SolVe Desktop application

## vSphere Distributed Switch (VDS) management tasks

For step-by-step instructions on managing the vSphere Distributed Switch (VDS) refer to the SolVe Desktop application or VMware vSphere documentation.

The following table lists the VDS administrative tasks by VxRail Manager software release.

Table 19 vSphere Distributed Switch (VDS) management tasks

Task	4.0.1xx	4.0.200	4.0.3xx	4.0.400	4.5.x	Notes
Add a new port group on default VDS	Yes	Yes	Yes	Yes	Yes	Refer to vSphere Documentation , "Add a Distributed Port Group"
Add a new port group on new VDS (Dell Node)	Yes	Yes	Yes	Yes	Yes	Refer to vSphere Documentation "Add a Distributed Port Group"
Add a new VDS(Dell Node)	Yes	Yes	Yes	Yes	Yes	Refer to vSphere Documentation , "Create a vSphere Distributed Switch"
Change default VDS NIOC configuration	Yes	Yes	Yes	Yes	Yes	Refer to the Solve Desktop application
Change IP of vMotion Network	Not supported	Yes	Yes	Yes	Yes	Contact Support
Change IP of vSAN Network (VxRail Supplied vCenter)	Not supported	Yes	Yes	Yes	Yes	Contact Support
Change IP of vSAN Network (Customer Supplied vCenter)	Not supported	Yes	Yes	Yes	Yes	Contact Support
Change Physical NIC Ports of Management Traffic	Not supported	Not supported	Not supported	Not supported	Not supported	
Change Physical NIC Ports of VM Network Traffic	Yes	Yes	Yes	Yes	Yes	Refer to the Solve Desktop application
Change Physical NIC Ports of vMotion Traffic	Yes	Yes	Yes	Yes	Not supported	Refer to the Solve Desktop application
Change Physical NIC Ports of vSAN Traffic	Not supported	Not supported	Not supported	Not supported	Not supported	
Change VLAN ID of Management, vCenter Server, and VxRail Management Network(VxRail Supplied vCenter)	Not supported	Yes	Yes	Yes	Yes	Refer to the Solve Desktop application
Change VLAN ID of VM Networks	Yes	Yes	Yes	Yes	Yes	Refer to the Solve Desktop application
Change VLAN ID of vMotion	Yes	Yes	Yes	Yes	Yes	Refer to the Solve Desktop application
Change VLAN ID of vSAN Network	Yes	Yes	Yes	Yes	Yes	Refer to the Solve Desktop application

Table 19 vSphere Distributed Switch (VDS) management tasks (continued)

Task	4.0.1xx	4.0.200	4.0.3xx	4.0.400	4.5.x	Notes
Enable Jumbo Frame	Yes	Yes	Yes	Yes	Yes	Refer to the Solve Desktop application
Enable network teaming LACP	Not supported	Not supported	Not supported	Not supported	Not supported	
Rename default port groups	Not supported	Yes	Yes	Yes	Yes	Refer to the Solve Desktop application
Rename default VDS	Not supported	Yes	Yes	Yes	Yes	Refer to the Solve Desktop application
Share network traffic with vSAN	Yes	Yes	Yes	Yes	Yes	Refer to the Solve Desktop application

## VM folder management tasks

For step-by-step instructions on managing the VM folder, refer to the Solve Desktop application.

The following table lists the VM folder related administrative tasks by VxRail Manager software release.

Table 20 VM folder management tasks

Task	4.0.1xx	4.0.200	4.0.3xx	4.0.400	4.5.x	Notes
Rename the VM folder	Yes	Yes	Yes	Yes	Yes	Refer to the Solve Desktop application

## vSphere Standard Switch (VSS) management tasks

For step-by-step instructions on managing the vSphere Standard Switch (VSS), refer to the VMware vSphere documentation.

The following table lists the VSS related administrative tasks by VxRail Manager software release.

Table 21 vSphere Standard Switch (VSS) management tasks

Task	4.0.1xx	4.0.200	4.0.3xx	4.0.400	4.5.x	Notes
Add a new VSS (Dell Node)	Yes	Yes	Yes	Yes	Yes	Refer to the vSphere Documentation, "Create a vSphere Standard Switch"



## VxRail Manager application management tasks

For step-by-step instructions on managing the VxRail Manager application, refer to the Solve Desktop application.

The following table lists the VxRail Manager application administrative tasks by VxRail Manager software release.

Table 22 VxRail Manager application management tasks

Task	4.0.1xx	4.0.200	4.0.3xx	4.0.400	4.5.x	Notes
Replace the VxRail Manager SSL Certificate	Not supported	Not supported	Yes	Yes	Yes	Refer to the Solve Desktop application
Change the VxRail Manager VM Hostname and IP Address	Not supported	Yes	Yes	Yes	Yes	Contact Support
Change the VxRail VM name	Not supported	Not supported	Yes	Yes	Yes	Refer to the Solve Desktop application
Change the VxRail Mystic account Password	Yes	Yes	Yes	Yes	Yes	Refer to the Solve Desktop application
Change the VxRail root account Password	Yes	Yes	Yes	Yes	Yes	Refer to the Solve Desktop application
Backup and restore key VxRail Manager files for recovery purposes	Not supported	Not supported	Not supported	Yes	Yes	Refer to the Solve Desktop application

Other administrative tasks

# CHAPTER 7

## External storage

Although each VxRail Appliance includes its own on-board storage resources, you can also use your appliance with external storage. Your VxRail Appliance supports NFS datastores and VMFS datastores over iSCSI. Using external storage doesn't affect VxRail features such as upgrades, reset, cluster shutdown and so on.

- NFS and VMFS over iSCSI datastores.....72
- Setting up external NFS storage for a VxRail Appliance..... 72
- Setting up external VMFS over iSCSI storage for a VxRail Appliance..... 73

## NFS and VMFS over iSCSI datastores

Your VxRail Appliance can be used with external storage with NFS datastores or VMFS over iSCSI datastores.

You may want to consider using external storage for workloads with high storage capacity requirements relative to the amount of required compute resources (CPU cores and RAM). External datastores allow you to scale VxRail Appliance compute resources separately from storage capacity. This can improved overall hardware usage levels. As a result it may be more cost-effective to use an external storage array such as Unity, VNX, Cloudarray, Isilon, or Windows/Linux storage services to provide additional storage for your VxRail Appliance cluster.

Note that external storage arrays must be managed separately. You cannot use the automated workflows in VxRail Manager for vSphere storage configuration and management with external datastores. Instead, you must use the vSphere Web Client to configure NFS and iSCSI on the ESXi hosts in the VxRail Appliance cluster, and to create NFS datastores and VMFS datastores over iSCSI on external arrays in the vCenter Server Appliance instance in the VxRail Appliance cluster. You must also use provisioning virtual machines on these external datastores in vCenter using the vSphere Web Client.

Refer to VMware vSphere documentation for details about using NFS and iSCSI storage with ESXi:

- Understanding Network File System Datastores
- About the Software iSCSI Adapter

## Setting up external NFS storage for a VxRail Appliance

Follow these guidelines to set up your VxRail Appliance for use with an NFS datastore.

### Before you begin

Many storage products provide NFS services, including including UNity, VNX, Isilon, CloudArray, and Windows/Linux NFS services. These guidelines assume that you have already installed and provisioned your VxRail Appliance cluster as well as your external NFS storage server.

The NFS datastore is created directly from vCenter on your VxRail Appliance cluster. A VLAN and a vSphere Distributed Switch (VDS) are used.

Refer to the VMware vSphere documentation, Set Up NFS Storage Environment, for details about NFS storage with ESXi.

### Procedure

1. Log in to the as an authorized user.
2. Click the vSphere Web Client home icon.
3. Click **Hosts and Clusters**.

In a VxRail Appliance cluster, NFS shared folders can be assigned to all connected nodes.

4. Click the cluster name in the list.
5. Click the **Related Objects** tab and then select **Datastores**.
6. Click the **Create a New Datastore** icon.

The **New Datastore** window is displayed.

7. In the **Location** window type the datastore name and select the location for the datastore.
8. In the **Type** window select **NFS**.
9. In the **Select NFS version** window choose an NFS version.
10. In the **Name and configuration** window, specify necessary information to add the external NFS storage.  
For **Folder** enter the Export Path of external storage array.
11. In the **Host accessibility** window, select the hosts in the VxRail Appliance cluster that will mount the datastore.  
You can select any or all hosts in the cluster.
12. In the **Ready to complete** window, review your entries and then click Finish to complete the datastore creation.

The new NFS datastore has been created.

#### Results

After the NFS datastore is created, you can:

- Provision VMs to the NFS datastore using the vSphere Web Client
- Perform standard operations such as Open Console, Power On, Shutdown VM, Suspend VM, and clone VM
- Migrate VMs between VxRail Appliance datastores and NFS datastores

NFS storage can take advantage of VMware Storage DRS features. In maintenance mode, VMs move to another node on the same NFS datastore automatically. The rebooted node is connected to the NFS datastore automatically.

## Setting up external VMFS over iSCSI storage for a VxRail Appliance

Follow these guidelines to set up your VxRail Appliance for use with a VMFS datastore over iSCSI.

Refer to the VMware documentation, *Configuring iSCSI Adapters and Storage*, while setting up external iSCSI storage for your VxRail Appliance.

The general steps to configure iSCSI storage for VxRail include the following :

1. Set up the iSCSI network portgroup.
2. Set up an active software iSCSI adapter.
3. Configure multipath policy for the iSCSI storage device.
4. Create a new VMFS datastore over iSCSI.

All the steps are done using the vSphere Web Client. A software iSCSI adapter allows standard NICs to connect ESXi hosts to a remote iSCSI target on the IP network. A VLAN and vSphere distributed switch are used. For details please refer to the VMware documentation, *About the Software iSCSI Adapter*.

### Set up the two iSCSI network portgroups

In the process of port binding, you create two virtual VMkernel adapters for each physical network. You then associate the VMkernel adapters with the software iSCSI

storage adapter created in the following steps. Create the distributed port groups first.

#### Note

The key is to avoid a network conflict with the VSAN. The uplink1/uplink2/vmnic0 description in the following procedure is suitable for most standard VxRail Appliance installations. Some VxRail Appliance models have extra Ethernet ports. You can use those ports as an alternate configuration. Replace uplink1 and uplink2 with the proper uplink numbers according to the extra ports.

#### Procedure

1. Log in to the vSphere Web Client as an authorized user.
2. In the vSphere Web Client, click the home icon and go to the **Networking** configuration page.
3. Highlight **VMware HCIA Distributed Switch** and then right-click it to create a new distributed port group.

The **New Distributed Port Group** window is displayed.

4. In the **Select name and location** window enter the port group name.  
A name containing "uplink1" is recommended.
5. In the **Configure settings** window enter the general properties and specify the related VLAN ID if you are using a VLAN.
6. In the **Ready to complete** window, review your entries and then click **Finish** to complete the port group creation.

The new distributed port group is created.

7. In the **Home** list, right-click the new portgroup and select **Edit Settings**.

The **Edit Settings** window is displayed.

8. If the port speed is 10Gb, change **Traffic Shaping** for the iSCSI portgroup as follows:
  - Set **Ingress traffic shaping: status** to **Enabled**, **Average bandwidth(kbit/s)** to **1048576**, **Peak bandwidth (kbit/s)** to **1048576**, and **Burst size (KB)** to **102400**.
  - Set **Egress traffic shaping: status** to **Enabled**, **Average bandwidth(kbit/s)** to **1048576**, **Peak bandwidth (kbit/s)** to **1048576**, and **Burst size (KB)** to **102400**.
9. Change **Teaming and failover** for the iSCSI portgroup as follows:
  - Keep **uplink1** as the **Active** uplink.
  - Move the other uplinks to **Unused**.
10. Create the second portgroup by repeating steps 3 through 9 with following value changes:
  - In step 4, a name containing "uplink2" is recommended.
  - in step 9, keep **uplink2** as the **Active** uplink.

#### Results

Set up of the two portgroups is complete.

## Create two VMkernel adapters on the hosts in the cluster

### Procedure

1. In the vSphere Web Client, click the home icon.
2. Select one of the hosts in the cluster.
3. Navigate to **Manage > Networking > VMkernel adapters**.
4. Click the **Add host networking** icon to add a VMkernel adapter with the first distributed port group you created.

The **Add Networking** window is displayed.

5. In the **Select connection type** window select **VMkernel Network Adapter**.
6. In the **Select target device** window choose **Select an existing network** and click **Browse**.
7. Select the "uplink1"-related distributed portgroup you created in the previous steps.
8. In the **IPv4 settings** window select **Use the static IPv4 setting**.  
This is defined according to your network environment. Make sure the IP address is unique in the network and accessible to the iSCSI target server IP address.
9. In the **Ready to complete** window, review your entries and then click **Finish** to complete the VMkernel creation.
10. Create the second VMkernel adapter by repeating steps 4 through 9 with following value change:
  - In step 7, select the "uplink2"-related distributed portgroup.

### Results

The two newly added VMkernel adapters are displayed and the detail info is shown in the lower panel. The network configuration for iSCSI on the host is finished.

## Set up an active software iSCSI adapter

### Before you begin

#### Note

The external iSCSI storage should already be prepared as required for your environment. This procedure does not cover the details of creating volumes used as external iSCSI storage.

### Procedure

1. In the left panel of the vSphere Web Client, click **vCenter > Hosts and Clusters** and select the ESXi host you previously configured.
2. Click **Manage > Storage > Storage Adapters** and click the **Add (+)** button to add a software iSCSI adapter.
3. In the **Add Software iSCSI Adapter** window, click **OK**.  
If needed, you can further configure the adapter in the **Adapter Details** section.  
The newly added storage adapter is displayed in the **Storage Adapters** list.

4. Click **Manage > Storage > Storage Adapters** and click **iSCSI Software Adapter**.

The created iSCSI name is listed in the **Identifier** column and **Properties** tab of the lower panel. Its format is `iqn.1998-01.com.vmware:xxxxxx-xxxxx`. You can change the iSCSI name by clicking **Edit** in the **Properties** tab of the lower panel.

5. Click **Network Port Binding** and click the **Add (+)** button.
6. In the **Bind .. with VMkernel Adapter** window, select the **VMkernel** network adapter that created in the previous steps.

The message "Due to recent configuration changes, a rescan of this storage adapter is recommended" is displayed.

7. Click the **Rescan the iSCSI Software Adapter** icon.
8. In the vSphere Web Client navigate to **Manage > Storage > Storage Adapters**.
9. Select the iSCSI software adapter you created and click **Targets > Dynamic Discovery**.
10. Click **Add**.
11. In the **Add Send Target Server** window specify the IP address of the iSCSI server and click **OK**.

The message "Due to recent configuration changes, a rescan of this storage adapter is recommended" is displayed.

---

#### Note

We recommend that you do not change the host iSCSI name after adding iSCSI storage information.

---

12. Click the **Rescan the iSCSI Software Adapter** icon.
- The warning message disappears and the iSCSI Target server is displayed in the lower panel.
13. Click **Static Discovery** to show the full target name of the iSCSI storage.

## Assign volumes on iSCSI storage

### Procedure

1. Repeat the procedures "Create a VMkernel adapter on the hosts in the cluster" and "Set up an active software iSCSI adapter" on the other nodes in the cluster.
2. Configure your iSCSI storage to assign storage volume mapping with the iSCSI name for all hosts.
3. In the vCenter Web Client choose each host and navigate to **Manage > Storage > Storage Adapters**.
4. Click the **Rescan the iSCSI Software Adapter** icon.

### Results

In the **Devices** tab, the new iSCSI volume is shown mapped to the ESXi host (iSCSI client).



## Configure multipath for the iSCSI volumes

To avoid any negative influence on VSAN performance, you must set multipath policy to **Fixed** and set **Preferred Path** on **vmnic0**. If you are using the extra ports on your VxRail Appliance (if equipped), set **Preferred Path** as the appropriate vmnic according to the extra ports.

### Procedure

1. In the vCenter Web Client choose each host and navigate to **Manage > Storage > Storage Device**.
2. Choose the iSCSI volume in the **Storage Devices** list.
3. In the **Properties** tab, click **Edit Multipathing....**  
The **Edit Multipathing Policies** window is displayed.
4. Select **Fixed(VMware)** for **Path selection policy** and then select the row in which **Runtime Name** contains **C0**.
5. Click **OK**.

### Results

In the **Properties** tab, the **Path Selection Policy** section shows **Fixed(VMware)** and **vmhbaxx:C0:Tx:Lx**.

## Mount a new VMFS datastore over iSCSI

To make full use of the new iSCSI storage, create a new datastore. This datastore can be mounted to all the hosts in the cluster.

### Procedure

1. In the vCenter Web Client select the cluster name.
2. Click **Related Objects > Datastores**, then click the **Add New Datastore** icon to create a new datastore over iSCSI.  
The **New Datastore** window is displayed.
3. In the **Type** window, select **VMFS**.
4. In the **Name and device selection** window specify the name for the datastore:
  - a. Choose one of the ESXi hosts for **Select a host to view its accessible disks/Luns**.
  - b. Select the available iSCSI volume in the lower table.
5. In the **Partition configuration** window, adjust the iSCSI datastore **Size**.  
It is recommended to set **Size** to **Use all available partitions**.
6. In the **Ready to complete** window, review your entries and then click **Finish** to complete the datastore creation.  
The datastore over iSCSI is listed. Its status is "Unknown".
7. Click the **Refresh** icon.  
The datastore status changes to "Normal".
8. In the **Storage** tab, check the connection status between the iSCSI datastore and the hosts.

### Results

After the VMFS over iSCSI datastore is created, you can:

- Provision VMs to the VMFS datastore using the vSphere Web Client
- Perform standard operations such as Open Console, Power On, Shutdown VM, Suspend VM, and clone VM
- Migrate VMs between VxRail Appliance datastores and iSCSI datastores
- Protect iSCSI external storage access with redundant path

In maintenance mode, VMs move to another node on the same iSCSI datastore automatically. The rebooted node is connected to the iSCSI datastore automatically. The iSCSI volume is kept as Remote Disk.

The benefit from the iSCSI multipath configuration is that if the physical adapter `vmnic0` is down, iSCSI IO switches to another `vmnic` without interruption. When `vmnic0` is back online, iSCSI IO will switch back to `vmnic0` without IO interruption.

# CHAPTER 8

## Replacing and adding hardware

You may be able to add or replace hardware components on your VxRail Appliance such as hard disk drives (HDDs), solid state drives (SSDs), power supplies, and so on. Only qualified personnel should perform these procedures. For some hardware components, you may need to contact Customer Support to arrange for repair or replacement.

Refer to the online support website or the SolVe Desktop application for hardware-specific information.

### **Using the SolVe Desktop application for VxRail Series hardware tasks**

Step-by-step hardware component tasks such as replacement and upgrade procedures are available through the SolVe Desktop application.

---

#### **▲ CAUTION**

**To avoid potential data loss, refer to the VxRail Series procedures in the SolVe Desktop application before performing any hardware replacement or upgrade procedures.**

---

To download the SolVe Desktop application, go to <https://support.emc.com> and click **SolVe** on the main page. Download and install the SolVe Desktop application on your computer.

You must have an online support account to use the SolVe Desktop application.

## Replacing and adding hardware

# Dell EMC VxRail Appliances on 14th Generation PowerEdge Servers

E Series Owner's Manual

## Notes, cautions, and warnings

 **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.

**Copyright © 2018 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.

# Contents

<b>1 Overview.....</b>	<b>5</b>
Front view.....	5
Left control panel view.....	6
Right control panel view.....	8
Back view of the appliance.....	9
Drive indicator codes.....	10
NIC indicator codes.....	11
Power supply unit indicator codes.....	11
iDRAC Direct LED indicator codes.....	14
Locating serial number of your appliance.....	14
Looking up your appliance serial number in VxRail Manager.....	15
Locating your physical VxRail Service Tag number.....	15
<b>2 Documentation matrix.....</b>	<b>16</b>
<b>3 Technical specifications.....</b>	<b>17</b>
Chassis dimensions.....	17
Chassis weight.....	18
Processor specifications.....	18
PSU specifications.....	18
Battery specifications.....	18
Expansion bus specifications.....	18
Memory specifications.....	19
Storage controller specifications.....	19
Hard drive specifications.....	19
Ports and connectors specifications.....	19
USB ports.....	19
NIC ports.....	20
Serial port.....	20
VGA ports.....	20
Internal dual SD Module.....	21
Environmental specifications.....	21
Particulate and gaseous contamination specifications.....	22
Standard operating temperature.....	23
Expanded operating temperature.....	23
Expanded operating temperature restrictions.....	24
Thermal restrictions.....	24
Ambient temperature limitations.....	24
<b>4 Initial setup and configuration.....</b>	<b>25</b>
<b>5 Pre-operating systems.....</b>	<b>26</b>

Options to manage the pre-operating system applications.....	26
iDRAC configuration.....	26
Log in to iDRAC.....	26
<b>6 Replacing and adding hardware.....</b>	<b>28</b>
Using Solve Desktop application for VxRail Series hardware tasks.....	28
Supported hardware components.....	28
System memory.....	29
General memory module installation guidelines.....	30
Expansion cards and expansion card riser.....	31
Expansion card installation guidelines.....	31
<b>7 Getting help.....</b>	<b>33</b>
Contacting Dell EMC.....	33
Registering for online support.....	33
Accessing support resources.....	33



# Overview

The VxRail™ E Series appliance is a hyper-converged appliance that supports:

- Up to two Intel Xeon Processor Scalable Family processors
- 10 x 2.5 inch hard drives or solid state drives (SSDs)
- 24 DIMM slots
- Two AC or DC redundant power supply units

**NOTE:** The appliance supports only internal, hot swappable hard drives.


**NOTE:** In this document, HDD generically refers to both HDD and SSD.

Topics:

- [Front view](#)
- [Back view of the appliance](#)
- [Drive indicator codes](#)
- [NIC indicator codes](#)
- [Power supply unit indicator codes](#)
- [iDRAC Direct LED indicator codes](#)
- [Locating serial number of your appliance](#)

## Front view


You can access the components such as power button, NMI button, appliance identification tag, appliance identification button, and USB ports from the front of the appliance. The diagnostic LEDs are prominently located on the front panel. The hot swappable hard drives are accessible from the front panel.



**Figure 1. Front view of 10 x 2.5 inch hard drive chassis**

**Table 1. Front panel features**

Item	Ports, panels, and slots	Icon	Description
1	Left control panel	N/A	Contains the appliance health and appliance ID, status LED. <ul style="list-style-type: none"> <li>• Status LED: Enables you to identify any failed hardware components. There are up to five status LEDs and an overall</li> </ul>

Item	Ports, panels, and slots	Icon	Description
			appliance health LED (Chassis health and appliance ID) bar. For more information, see the <a href="#">Status LED indicators</a> section.
2	VGA port		Enables you to connect a display device to the appliance. For more information, see the <a href="#">Technical specifications</a> section.
3	Right control panel	N/A	Contains the power button, USB port, iDRAC Direct micro-USB, and the iDRAC Direct status LED.
4	Drive slots	N/A	Enable you to install drives that are supported on your appliance. For more information about drives, see the <a href="#">Technical specifications</a> section.

## Left control panel view





Figure 2. Left control panel view


Table 2. Left control panel features





Item	Indicator	Icon	Description
1	Status LED indicators	N/A	Indicate the status of the appliance. For more information, see the <a href="#">Status LED indicators</a> section.
2	Appliance health and appliance ID indicator		Indicates the appliance health. For more information, see the appliance health and appliance ID indicator codes section.

## Status LED indicators

**NOTE:** The status LED indicators are always off and only turns on to a solid amber if any error occurs.

Table 3. Status LED indicators and descriptions

Icon	Description	Condition	Corrective action
	Drive indicator	The indicator turns solid amber if there is a drive error.	<ul style="list-style-type: none"> <li>Check the System Event Log to determine if the drive has an error.</li> <li>Run the appropriate Online Diagnostics test. Restart the appliance and run embedded diagnostics (ePSA).</li> </ul>

Icon	Description	Condition	Corrective action
	Temperature indicator	The indicator turns solid amber if the appliance experiences a thermal error (for example, the ambient temperature is out of range or there is a fan failure).	<ul style="list-style-type: none"> <li>If the drives are configured in a RAID array, restart the appliance, and enter the host adapter configuration utility program.</li> </ul> <p>Ensure that none of the following conditions exist:</p> <ul style="list-style-type: none"> <li>A cooling fan has been removed or has failed.</li> <li>Appliance cover, air shroud, memory module blank, or back filler bracket is removed.</li> <li>Ambient temperature is too high.</li> <li>External airflow is obstructed.</li> </ul> <p>If the problem persists, see the <a href="#">Getting help</a> section.</p>
	Electrical indicator	The indicator turns solid amber if the appliance 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 appliance 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 <a href="#">Getting help</a> section.
	Memory indicator	The indicator turns solid amber if a memory error occurs.	Check the System Event Log or appliance messages for the location of the failed memory. Reseat the memory module. If the problem persists, see the <a href="#">Getting help</a> section.
	PCIe indicator	The indicator turns solid amber if a PCIe card experiences an error.	Restart the appliance. Update any required drivers for the PCIe card. Reinstall the card. If the problem persists, see the <a href="#">Getting help</a> section.
			<p><b>NOTE:</b> For more information about the supported PCIe cards, see the <a href="#">Expansion card installation guidelines</a> section.</p>

## Appliance health and appliance ID indicator codes

The appliance health and appliance ID indicator is located on the left control panel of your appliance.



Figure 3. Appliance health and appliance ID indicators

Table 4. Appliance health and appliance ID indicator codes

Appliance health and appliance ID indicator code	Condition
Solid blue	Indicates that the appliance is turned on, appliance is healthy, and appliance ID mode is not active. Press the appliance health and appliance ID button to switch to appliance ID mode.
Blinking blue	Indicates that the appliance ID mode is active. Press the appliance health and appliance ID button to switch to appliance health mode.
Solid amber	Indicates that the appliance is in fail-safe mode. If the problem persists, see the <a href="#">Getting help</a> section.

## Appliance health and appliance ID indicator code

Blinking amber

## Condition

Indicates that the appliance is experiencing a fault. Check the Appliance Event Log or the LCD panel, if available on the bezel, for specific error message. For more information about error messages, see the *Dell Event and Error Messages Reference Guide* at [Dell.com/openmanagemanuals](https://Dell.com/openmanagemanuals) > **OpenManage software**.

## Right control panel view








Figure 4. Right control panel view

Table 5. Right control panel features

Item	Indicator, button, or connector	Icon	Description
1	Power button		Indicates if the appliance is turned on or off. Press the power button to manually turn on or off the appliance. <b>NOTE: Press the power button to gracefully shut down an ACPI-compliant operating system.</b>
2	USB port (2)		The USB ports are 4-pin, 2.0-compliant. These ports enable you to connect USB devices to the appliance.
3	iDRAC Direct port		The iDRAC Direct port is micro USB 2.0-compliant. This port enables you to access the iDRAC Direct features. For more information, see the iDRAC User's Guide at <a href="https://Dell.com/idracmanuals">Dell.com/idracmanuals</a> .
4	iDRAC Direct LED	N/A	The iDRAC Direct LED indicator lights up to indicate that the iDRAC Direct port is connected. For more information, see the iDRAC Direct LED indicator codes section.
5	VGA port		Enables you to connect a display device to the appliance. For more information, see the <a href="#">Technical specifications</a> section.

# Back view of the appliance

You can access the components such as appliance identification button, power supply sockets, cable management arm connectors, iDRAC storage media, NIC ports, and USB ports from the back of the appliance. Most of the expansion card ports can be accessed from the back panel. The hot swappable power supply units are accessible from the back panel.










Figure 5. Back view of appliance with 3 PCIe expansion slots

Table 6. 10 X 2.5 inch drive appliance with 3 PCIe expansion slot

Item	Ports, panels, or slots	Icon	Description
1	PCIe expansion card slot(s)	N/A	The expansion slot(s) enable you to connect PCI Express expansion cards. For more information on the expansion cards that are supported on your appliance, see the <a href="#">Expansion card guidelines</a> section.
2	Power supply unit (2)	N/A	For more information about the PSU configurations, see the <a href="#">Technical Specifications</a> section.
3	NIC port (4)		The NIC ports that are integrated on the network daughter card (NDC) provide network connectivity. For more information about the supported configurations, see the <a href="#">Technical specifications</a> section.
4	USB 3.0 (2)		The USB ports are 9-pin and 3.0-compliant. These ports enable you to connect USB devices to the appliance.
5	VGA port		Enables you to connect a display device to the appliance. For more information, see the <a href="#">Technical Specifications</a> section.
6	Serial port		Enables you to connect a serial device to the appliance. For more information, see the <a href="#">Technical Specifications</a> section.
7	iDRAC9 Enterprise port		Enables you to remotely access iDRAC. For more information, see the <a href="#">iDRAC User's Guide</a> at <a href="http://Dell.com/idracmanuals">Dell.com/idracmanuals</a> .
8	CMA power port	N/A	The Cable Management Arm (CMA) power port enables you to connect to the CMA.
9	Appliance identification button		The Appliance Identification (ID) button is available on the front and back of the appliances. Press the button to identify a appliance in a rack by turning on the appliance ID button. You can also use the appliance ID button to reset iDRAC and to access BIOS using the step through mode.

# Drive indicator codes

Each hard drive carrier has an activity indicator and a status indicator. The indicators provide information about the current status of the hard drive. The activity LED indicates whether the hard drive is currently in use or not. The status LED indicates the power condition of the hard drive.




Figure 6. Drive indicators


- 1 drive activity LED indicator
- 2 drive status LED indicator
- 3 drive

Table 7. Drive indicator codes

Drive status indicator code	Condition
Flashes green twice per second	Identifying drive or preparing for removal.
Off	Drive ready for removal.
	<b>NOTE:</b> The drive status indicator remains off until all drives are initialized after the appliance is turned on. Drives are not ready for removal during this time.
Flashes green, amber, and then turns off	Predicted drive failure.
Flashes amber four times per second	Drive failed.
Solid green	Drive online.

# NIC indicator codes

The NIC on the back panel has an indicator that provides information about the network activity and link status. The activity LED indicates whether the NIC is currently connected or not. The link LED indicates the speed of the connected network.



**Figure 7. NIC indicators**


- 1 link indicator
- 2 activity indicator

**Table 8. NIC indicators**

Convention	Status	Condition
A	Link and activity indicators are off	The NIC is not connected to the network.
B	Link indicator is green	The NIC is connected to a valid network at its maximum port speed (1 Gbps or 10 Gbps).
C	Link indicator is amber	The NIC is connected to a valid network at less than its maximum port speed.
D	Activity indicator is flashing green	Network data is being sent or received.

# Power supply unit indicator codes

AC power supply units (PSUs) have an illuminated translucent handle that serves as an indicator and DC PSUs have an LED that serves as an indicator. The indicator shows whether power is present or a power fault has occurred.



**Figure 8. AC PSU status indicator**

1 AC PSU status indicator/handle

**Table 9. AC PSU status indicators**

Convention	Power indicator pattern	Condition
A	Green	A valid power source is connected to the PSU and the PSU is operational.
B	Flashing green	When the firmware of the PSU is being updated, the PSU handle flashes green.
C	Flashing green and turns off	When hot-adding a PSU, the PSU handle flashes green five times at 4 Hz rate and turns off. This indicates a PSU mismatch with respect to efficiency, feature set, health status, and supported voltage.
		<p><b>NOTE:</b> Ensure that both the PSUs are of the same capacity.</p> <p><b>CAUTION:</b> For AC PSUs, use only PSUs with the Extended Power Performance (EPP) label on the back.</p>
D	Flashing amber	Indicates a problem with the PSU.
		<p><b>CAUTION:</b> When correcting a PSU mismatch, replace only the PSU with the flashing indicator. Swapping the PSU to make a matched pair can result in an error condition and unexpected appliance shutdown. To change from a high output configuration to a low output configuration or vice versa, you must turn off the appliance.</p> <p><b>CAUTION:</b> AC PSUs support both 220 V and 110 V input voltages. When two identical PSUs receive different input voltages, they can output different wattages, and trigger a mismatch.</p> <p><b>CAUTION:</b> If two PSUs are used, they must be of the same type and have the same maximum output power.</p> <p><b>CAUTION:</b> Combining AC and DC PSUs is not supported and triggers a mismatch.</p>



Convention	Power indicator pattern	Condition
E	Not lit	Power is not connected.




Figure 9. DC PSU status indicator

1 DC PSU status indicator


Table 10. DC PSU status indicators

Convention	Power indicator pattern	Condition
A	Green	A valid power source is connected to the PSU and that the PSU is operational.
B	Flashing green	When hot-adding a PSU, the PSU indicator flashes green. This indicates that there is a PSU mismatch with respect to efficiency, feature set, health status, and supported voltage. Ensure that both the PSUs are of the same capacity.
C	Flashing amber	Indicates a problem with the PSU. <b>CAUTION:</b> When correcting a PSU mismatch, replace only the PSU with the flashing indicator. Swapping the PSU to make a matched pair can result in an error condition and unexpected appliance shutdown. To change from a High Output configuration to a Low Output configuration or vice versa, you must turn off the appliance. <b>CAUTION:</b> If two PSU are used, they must be of the same type and have the same maximum output power. <b>CAUTION:</b> Combining AC and DC PSU is not supported and triggers a mismatch.
D	Not lit	Power is not connected.

# iDRAC Direct LED indicator codes

The iDRAC Direct LED indicator lights up to indicate that the port is connected and is being used as a part of the iDRAC subsystem.

**NOTE:** The iDRAC Direct LED indicator does not turn on when the USB port is used in the USB mode.



**Figure 10. iDRAC Direct LED indicator**

1 iDRAC Direct status indicator

The following table describes iDRAC Direct activity when configuring iDRAC Direct by using the management port (USB XML Import).

**Table 11. iDRAC Direct LED indicators**

Convention	iDRAC Direct LED indicator pattern	Condition
A	Green	Turns green for a minimum of two seconds to indicate the start and end of a file transfer.
B	Flashing green	Indicates file transfer or any operation tasks.
C	Green and turns off	Indicates that the file transfer is complete.
D	Not lit	Indicates that the USB is ready to be removed or that a task is complete.

The following table describes iDRAC Direct activity when configuring iDRAC Direct by using your laptop and cable (Laptop Connect):

**Table 12. iDRAC Direct LED indicator patterns**

iDRAC Direct LED indicator pattern	Condition
Solid green for two seconds	Indicates that the laptop is connected.
Flashing green (on for two seconds and off for two seconds)	Indicates that the laptop connected is recognized.
Turns off	Indicates that the laptop is unplugged.

## Locating serial number of your appliance

To get support for your appliance, use the VxRail Appliance serial number, also called the Product Serial Number Tag (PSNT). The PSNT is a 14-digit number used to identify your appliance to Dell EMC support.

**NOTE:** Only use the VxRail Appliance serial number to contact Customer Support. Sometimes, you may need to supply the 7-digit Service Tag number.

There are two identification tags on your appliance:

- The VxRail appliance serial number—You can find the serial number in **VxRail Manager**, or printed on the information tag.
- The Service Tag — You can find the Service Tag printed on the physical appliance.

## Looking up your appliance serial number in VxRail Manager

- 1 In **VxRail Manager**, on the left navigation bar, click **Health**.
- 2 To display appliance information, click **Physical**.
- 3 Observe the VxRail Appliance serial number, listed under the appliance ID as the **PSNT**.

## Locating your physical VxRail Service Tag number

Your hardware is identified by a unique Service Tag number. The Service Tag is found on the front of the appliance by pulling out the information tag.

Alternatively, the information may be on a sticker on the chassis of the appliance. This information is used by Dell EMC to route support calls to the appropriate personnel.

**NOTE:** Use the 7-digit Service Tag, only if instructed by Technical Support.




Figure 11. VxRail Service Tag

# Documentation matrix

The documentation matrix provides information on documents that you can refer while setting up and managing your appliance.

**Table 13. Documentation matrix**

Document	Provides information about...	Location
<b>Software Documents</b>		
Online help in the VxRail Manager UI	all admin tasks, licensing, and product architecture information.	VxRail Manager Online Help
Administrator Guide	admin tasks and conceptual information.	<a href="http://emc.com/vxrailsupport">emc.com/vxrailsupport</a>
Release Notes	the product and any critical information about the release.	<a href="http://emc.com/vxrailsupport">emc.com/vxrailsupport</a>
<b>Hardware Documents</b>		
Getting started with your appliance	setting up your appliance.	<a href="http://emc.com/vxrailsupport">emc.com/vxrailsupport</a>
Owner's Manual	hardware details for your appliance along with technical specifications.	<a href="http://emc.com/vxrailsupport">emc.com/vxrailsupport</a>
<b>Service Procedure Documents</b>		
Solve Desktop application	the Solve Desktop application. The Solve Desktop application gathers critical information from EMC product guides and combines it with expert Dell EMC support advice to generate a procedure document that is concise and task driven.	EMC Online Support site



**NOTE:** Download the Solve Desktop application, all generators are available within the Solve Desktop.

## Technical specifications

The technical and environmental specifications of your appliance are outlined in this section.

### Chassis dimensions




Figure 12. Chassis dimensions

Table 14. Dimension details

Appliance	Xa	Xb	Y	Za (with bezel)	Za (without bezel)	Zb*	Zc
10 x 2.5 inches	482.0 mm (18.97 inches)	434.0 mm (17.08 inches)	42.8 mm (1.68 inches)	35.84 mm (1.41 inches)	22.0 mm (0.87 inches)	733.82 mm (29.61 inches)	772.67 mm (30.42 inches)

# Chassis weight

Table 15. Chassis weight

System	Maximum weight (with all hard drives/SSDs)
VxRail E Series	21.9 kg (48.28 lbs)

# Processor specifications

The VxRail E Series appliance supports up to two Intel Xeon Processor Scalable Family processors.

# PSU specifications

The VxRail E Series appliance supports two AC or DC power supply units (PSUs).

Table 16. PSU specifications

PSU	Class	Heat dissipation (maximum)	Frequency	Voltage
1100 W DC	Gold	4416 BTU/hr	50/60 Hz	-(48-60) V DC
1100 W Mixed Mode HVDC (for China and Japan only)	Platinum	4100 BTU/hr	50/60 Hz	100-240 V AC and 200-380 V DC
1100 W AC	Platinum	4100 BTU/hr	50/60 Hz	100-240 V AC, autoranging

**NOTE:** If an appliance with 1100 W AC or HVDC PSU operates from 100 to 120V, the power rating per PSU is derated to 1050 W.

**NOTE:** Heat dissipation is calculated using the PSU wattage rating.

**NOTE:** This appliance is also designed to connect to the IT power systems with a phase to phase voltage not exceeding 240 V.

# Battery specifications

The VxRail E Series appliance supports CR 2032 3.0-V lithium coin cell battery.

# Expansion bus specifications

The VxRail E Series appliance supports PCI express (PCIe) generation 3 expansion cards, which must be available on the system board by using expansion card risers.

Table 17. Expansion card riser specifications

Supported VxRail configurations	Riser configuration and supported risers	Slot description	PCIe slots (Height and length)	Processor connection
VxRail E Series appliance	Riser configuration 2 (1A +2A)	Three x16 slots	Slot 1: x16 low profile, half length Slot 2: x16 low profile, half length	Processor 1 Processor 1

Supported VxRail configurations	Riser configuration and supported risers	Slot description	PCIe slots (Height and length)	Processor connection
			Slot 3: x16 low profile, half length	Processor 2
	Riser configuration 4 (1A)	Two x16 slots	Slot 2: x16 low profile, half length	Processor 1

## Memory specifications

Table 18. Memory specifications

DIMM type	DIMM rank	DIMM capacity	Single processor		Dual processors	
			Minimum RAM	Maximum RAM	Minimum RAM	Maximum RAM
LRDIMM	Quad rank	64 GB	384 GB	768 GB	768 GB	1536 GB
RDIMM	Dual rank	16 GB	96 GB	192 GB	192 GB	384 GB
RDIMM	Dual rank	32 GB	192 GB	384 GB	384 GB	768 GB

## Storage controller specifications

The VxRail E Series appliance supports:

- Internal storage controller cards: HBA330 and Boot Optimized Server Storage (BOSS-S1).

## Hard drive specifications

The VxRail E Series appliance supports NVMe, SAS, SATA, Nearline SAS hard drives or SSDs.

Table 19. Supported drive options

Ten drive appliances

Up to ten 2.5 inch, hot swappable NVMe, SAS, SAS/SATA SSD, or Nearline SAS hard drives.

**NOTE: VxRail E Series appliance supports NVMe drives in cache tier, slots 8 and 9 only.**

## Ports and connectors specifications

The VxRail E Series appliance supports USB ports, NIC ports, VGA ports, serial connector, and an IDSDM card that supports two internal dual SD module.

### USB ports

The following table provides more information about the USB specifications:

**Table 20. USB specifications**

Appliance	Front panel	Back panel	Internal
Ten hard drive appliances	Two 4-pin, USB 2.0-compliant port  One 5-pin micro USB 2.0 management port	Two 9-pin, USB 3.0-compliant ports  N/A	One 9-pin, USB 3.0-compliant ports  N/A

**NOTE:** The micro USB 2.0-compliant port on the front panel can only be used as an iDRAC Direct or a management port.

## NIC ports

The VxRail E Series appliance supports four Network Interface Controller (NIC) ports on the back panel, which is available in one of the following NIC configurations:

- Four RJ-45 ports that support 100 M, 1 G and 10 Gbps
- Four SFP+ ports that support 10 Gbps

## Serial port

The serial connector connects a serial device to the appliance. The VxRail E Series appliance supports one serial connector on the back panel, which is a 9-pin connector, Data Terminal Equipment (DTE), 16550-compliant.

## VGA ports

The Video Graphic Array (VGA) port enables you to connect the appliance to a VGA display. The VxRail E Series appliance supports one 15-pin VGA port on the front and back of appliance.

## Video specifications

The VxRail E Series appliance supports integrated VGA controller with 4 MB SPI capacity.

**Table 21. Supported video resolution options**

Resolution	Refresh rate (Hz)	Color depth (bits)
640 x 480	60, 70	8, 16, 32
800 x 600	60, 75, 85	8, 16, 32
1024 x 768	60, 75, 85	8, 16, 32
1152 x 864	60, 75, 85	8, 16, 32
1280 x 1024	60, 75	8, 16, 32
1440 x 900	60	8, 16, 32
1920 x 1200	60	8, 16, 32



## Internal dual SD Module

The Internal Dual SD module (IDSDM) provides a redundant SD card solution. The IDSDM card offers the following feature:

- Dual card operation — Maintains a mirrored configuration by using SD cards in both the slots and provides redundancy.

**NOTE:** The IDSDM on the VxRail appliance is pre-configured for appliance bare metal recovery. It is recommended not to change any settings.

## Environmental specifications

**NOTE:** For additional information about environmental measurements for specific appliance configurations, see [Dell.com/environmental\\_datasheets](http://Dell.com/environmental_datasheets).

**Table 22. Temperature specifications**

Temperature	Specifications
Storage	-40°C to 65°C (-40°F to 149°F)
Continuous operation (for altitude less than 950 m or 3117 ft)	10°C to 35°C (50°F to 95°F) with no direct sunlight on the equipment. <b>NOTE:</b> Maximum of 205 W, 28 core processor is supported in appliances. <b>NOTE:</b> Certain configurations may have ambient temperature restrictions. For more information see the Ambient temperature limitations section.
Fresh air	For information about fresh air, see <a href="#">Expanded Operating Temperature</a> section.
Maximum temperature gradient (operating and storage)	20°C/h (68°F/h)

**Table 23. Relative humidity specifications**

Relative humidity	Specifications
Storage	5% to 95% RH with 33°C (91°F) maximum dew point. Atmosphere must be non-condensing at all times.
Operating	10% to 80% relative humidity with 29°C (84.2°F) maximum dew point.

**Table 24. Maximum vibration specifications**

Maximum vibration	Specifications
Operating	0.26 G <sub>rms</sub> at 5 Hz to 350 Hz (all operation orientations).
Storage	1.88 G <sub>rms</sub> at 10 Hz to 500 Hz for 15 min (all six sides tested).

**Table 25. Maximum shock specifications**

Maximum shock	Specifications
Operating	Six consecutively executed shock pulses in the positive and negative x, y, and z axes of 6 G for up to 11 ms.
Storage	Six consecutively executed shock pulses in the positive and negative x, y, and z axes (one pulse on each side of the appliance) of 71 G for up to 2 ms.

**Table 26. Maximum altitude specifications**

Maximum altitude	Specifications
Operating	3048 m (10,000 ft)
Storage	12,000 m (39,370 ft)

**Table 27. Operating temperature de-rating specifications**

Operating temperature de-rating	Specifications
Up to 35°C (95°F)	Maximum temperature is reduced by 1°C/300 m (1°F/547 ft) above 950 m (3,117 ft).
35°C to 40°C (95°F to 104°F)	Maximum temperature is reduced by 1°C/175 m (1°F/319 ft) above 950 m (3,117 ft).
40°C to 45°C (104°F to 113°F)	Maximum temperature is reduced by 1°C/125 m (1°F/228 ft) above 950 m (3,117 ft).

## Particulate and gaseous contamination specifications

The following table defines the limitations that help avoid any equipment damage or failure from particulates and gaseous contamination. If the levels of particulates 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 28. Particulate 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.  <b>NOTE:</b> This condition applies only to data center environments. Air filtration requirements do not apply to IT equipment designed to be used outside a data center, in environments such as an office or factory floor.  <b>NOTE:</b> 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.  <b>NOTE:</b> This condition applies to data center and non-data center environments.

## Particulate contamination

Corrosive dust

## Specifications

- Air must be free of corrosive dust.
- Residual dust present in the air must have a deliquescent point less than 60% relative humidity.

**NOTE:** This condition applies to data center and non-data center environments.

Table 29. Gaseous contamination specifications

## Gaseous contamination

Copper coupon corrosion rate

Silver coupon corrosion rate

## Specifications

<300 Å/month per Class G1 as defined by ANSI/ISA71.04-1985.

<200 Å/month as defined by AHSRAE TC9.9.

**NOTE:** Maximum corrosive contaminant levels measured at ≤50% relative humidity.

# Standard operating temperature

The standard operating temperature for altitude less than 950 meters or 3117 feet ranges from 10°C to 35°C (50°F to 95°F) with no direct sunlight on the equipment.

# Expanded operating temperature

Table 30. Expanded operating temperature specifications

## Expanded operating temperature

Continuous operation

≤ 1% of annual operating hours

## Specifications

5°C to 40°C at 5% to 85% RH with 29°C dew point.

**NOTE:** Outside the standard operating temperature (10°C to 35°C), the appliance can operate continuously in temperatures as low as 5°C and as high as 40°C.

For temperatures between 35°C and 40°C, de-rate maximum allowable temperature by 1°C per 175 m above 950 m (1°F per 319 ft).

-5°C to 45°C at 5% to 90% RH with 29°C dew point.

**NOTE:** Outside the standard operating temperature (10°C to 35°C), the appliance can operate down to -5°C or up to 45°C for a maximum of 1% of its annual operating hours.

For temperatures between 40°C and 45°C, de-rate maximum allowable temperature by 1°C per 125 m above 950 m (1°F per 228 ft).

**NOTE:** When operating in the expanded temperature range:

- The performance of the appliance may be impacted.
- The ambient temperature warnings may be reported in the System Event Log.

## Expanded operating temperature restrictions

- Do not perform a cold startup below 5°C.
- The operating temperature specified is for a maximum altitude of 3050 m (10,000 ft).
- Redundant power supply units are required.
- Non-Dell EMC qualified peripheral cards and/or peripheral cards greater than 25 W are not supported.

## Thermal restrictions

The following table lists the configurations required for efficient cooling.

**Table 31. Thermal restrictions configuration**

Configuration	Number of processors	Heatsink	Processor/ DIMM blank	DIMM blanks	Maximum number of DIMM blanks	Fan
2.5 inch hard drives x 10	1	One 1U standard heat sink for CPU ≤ 165 W	Not required	Required for processor 1	11 blanks	Five standard fans
		One 1U 2-pipe heat sink for CPU=200/205 W	Required			Eight high performance fans
	2	Two 1U standard heat sink for CPU ≤ 165 W	Not required			Eight standard fans
		Two 1U 2-pipe heat sink for CPU=200/205 W		Required	22 blanks	Eight high performance fans

## Ambient temperature limitations

The following table lists configurations that require ambient temperature less than 35°C.

**NOTE:** The ambient temperature limit must be adhered to ensure proper cooling and to avoid excess processor throttling, which may impact appliance performance.

**Table 32. Configuration based ambient temperature restrictions**

Appliance	Front Backplane	Processor Thermal Design Power	Processor Heat Sink	Fan Type	Ambient Restriction
VxRail E Series Appliances	10 x 2.5 inch SAS/ SATA hard drives	200 W, 205 W	2 pipe 1U high performance	High performance fan	30°C

# Initial setup and configuration

For assistance on installation and deployment services, contact your Dell EMC account team or your reseller for installation services.

**⚠ WARNING:** During the VxRail deployment process, an iDRAC account named vxadmin or PTAdmin is created. This account provides hardware information to the VxRail Manager and is required for the VxRail Manager and the cluster to function properly. Ensure that you do not modify or delete this account.

**📌 NOTE:** Do not install the appliance into a rack, or turn on the appliance during the deployment process.

# Pre-operating systems

You can manage basic settings and features of a appliance without booting to the operating system by using the appliance firmware.

## NOTE:

- This appliance requires installation and deployment services. Do not rack this appliance, or turn on the power. Contact your Dell EMC Account Team or your reseller for setting up your appliance.
- Dell EMC has optimized your appliance. Do not change any of these settings.

Topics:

- [Options to manage the pre-operating system applications](#)
- [iDRAC configuration](#)

## Options to manage the pre-operating system applications

Your appliance has the following options to manage the pre-operating system applications:

- System Setup
- Boot Manager
- Lifecycle Controller
- Preboot Execution Environment (PXE)

**NOTE:** Dell EMC has optimized your appliance and recommends that you do not change any of these settings.

## iDRAC configuration

The Integrated Dell Remote Access Controller (iDRAC) is designed to make appliance administrators more productive and improve the overall availability of Dell EMC appliances. iDRAC alerts administrators to appliance issues, helps them perform remote appliance management, and reduces the need for physical access to the appliance.

## Log in to iDRAC

You can log in to iDRAC as:

- iDRAC user
- Microsoft Active Directory user
- Lightweight Directory Access Protocol (LDAP) user

The default user name and password are `root` and `calvin`. You can also log in by using Single Sign-On or Smart Card.

**NOTE:** You must have iDRAC credentials to log in to iDRAC.

For more information about logging in to iDRAC and iDRAC licenses, see the latest *Integrated Dell Remote Access Controller User's Guide* at [Dell.com/idracmanuals](https://www.dell.com/idracmanuals).

The iDRAC IP address is pre-configured for DHCP. This can be changed to a static IP address by logging into iDRAC.

**NOTE:**

- To access iDRAC, connect the network cable to the Ethernet connector 1 on the system board.
- Ensure that you change the default user name and password after setting up the iDRAC IP address.

## Replacing and adding hardware

Adding or replacing hardware component procedures on your VxRail E Series appliance, such as hard disk drives (HDDs), solid state drives (SSDs), and power supply units must be performed only by Dell EMC certified service technicians. For certain hardware components, you may need to contact Customer Support for repair or replacement.

### Using SolVe Desktop application for VxRail Series hardware tasks

Step-by-step hardware component tasks such as replacement and upgrade procedures are available through the SolVe Desktop application.

You must have an online support account to use the SolVe Desktop application.

**⚠ WARNING:** The VxRail Series procedures in the SolVe Desktop application for replacing hardware or any upgrade procedures must be performed only by Dell EMC certified service technicians.

**⚠ CAUTION:** To avoid data loss, ensure that you refer to the VxRail Series procedures in the SolVe Desktop application before replacing hardware or performing any upgrade procedures.

- 1 Log in to the EMC Online Support site.
- 2 Click **Solve** on the main page.
- 3 Click the download link for the SolVe Desktop application.
- 4 Save the executable file and then run it to install the SolVe Desktop.

### Supported hardware components

Following table describes the supported hardware components for E series appliance.

**Table 33. Supported hardware components**

Hardware component	Customer Replaceable Unit (CRU)	Field Replaceable Unit (FRU)
System Memory	Yes	No
Hard Drive	Yes	No
Solid State Drive (SAS/SATA)	Yes	No
Solid State Drive (NVMe)	Yes	No
PCIe Network Interface Cards	Yes	No
Micro SDHC Card	Yes	No
Power Supply Unit	Yes	No
Processors	No	Yes
System Motherboard	No	Yes
Host Bus Adapter (HBA330)	No	Yes
BOSS controller card and M.2 SATA disk	No	Yes
Network Daughter Card (NDC)	No	Yes



**NOTE:** The preceding table provides a non-exhaustive list of FRUs that reflects the common top-level assembly parts.

## System memory

The appliance supports DDR4 registered DIMMs (RDIMMs) and load reduced DIMMs (LRDIMMs). Memory holds the instructions that are executed by the processor.


**CAUTION:** To avoid data loss, ensure that you refer to the procedures in the SolVe Desktop application before performing any memory or expansion card replacement or upgrade procedures.

**NOTE:** MT/s indicates DIMM speed in MegaTransfers per second.

Memory bus operating frequency can be 2666 MT/s, 2400 MT/s, or 2133 MT/s depending on the following factors:

- DIMM type (RDIMM or LRDIMM)
- Number of DIMMs populated per channel
- System profile selected (for example, Performance Optimized, Custom, or Dense Configuration Optimized)
- Maximum supported DIMM frequency of the processors

Your appliance contains 24 memory sockets split into two sets of 12 sockets, one set per processor. Each 12-socket set is organized into six channels. In each channel, the release tabs of the first socket are marked white and the second socket are marked black.



**Figure 13. Memory socket locations**

Memory channels are organized as follows:

**Table 34. Memory channels**

Processor	Channel 0	Channel 1	Channel 2	Channel 3	Channel 4	Channel 5
Processor 1	Slots A1 and A7	Slots A2 and A8	Slots A3 and A9	Slots A4 and A10	Slots A5 and A11	Slots A6 and A12
Processor 2	Slots B1 and B7	Slots B2 and B8	Slots B3 and B9	Slots B4 and B10	Slots B5 and B11	Slots B6 and B12

## General memory module installation guidelines

**NOTE:** Memory configurations that fail to observe these guidelines can prevent your appliance from booting, stop responding during memory configuration, or operating with reduced memory.

The following are the recommended guidelines for installing memory modules:

- RDIMMs and LRDIMMs must not be mixed.
- If memory modules with different speeds are installed, they will operate at the speed of the slowest installed memory module(s) or slower depending on appliance DIMM configuration.
- Populate memory module sockets only if a processor is installed. For single-processor appliance, sockets A1 to A12 are available. For dual-processor appliance, sockets A1 to A12 and sockets B1 to B12 are available.
- Populate all the sockets with white release tabs first, and then the black release tabs.
- Mixing of more than one memory module capacities in the appliance is not supported.
- In a dual-processor configuration, the memory configuration for each processor should be identical. For example, if you populate socket A1 for processor 1, then populate socket B1 for processor 2, and so on.
- VxRail Appliances require that you populate six memory modules per processor for single processor configurations and twelve memory modules per processor for dual-processor configurations (one DIMM per channel) at a time to maximize performance.

## Expansion cards and expansion card riser

An expansion card in the appliance is an add-on card that can be inserted into an expansion slot on the system board or riser card to add enhanced functionality to the appliance through the expansion bus.

- ① **NOTE:** To avoid data loss, ensure that you refer to the procedures in the Solve Desktop application before performing any memory or expansion card replacement or upgrade procedures.
- ① **NOTE:** A System Event Log (SEL) event is logged if an expansion card riser is unsupported or missing. It does not prevent your appliance from turning on and no BIOS POST message or F1/F2 pause is displayed.

## Expansion card installation guidelines

Your appliance supports PCI Express Generation 3 expansion cards.

**Table 35. Expansion card riser configurations**

Expansion card riser	PCIe slots on the riser	Processor connection	Height	Length	Slot width
Riser 1A	Slot 1	Processor 1	Low Profile	Half Length	x16
	Slot 2	Processor 1	Low Profile	Half Length	x16
Riser 2A	Slot 3	Processor 2	Low Profile	Half Length	x16

The following table provides guidelines for expansion cards priority to ensure proper cooling and mechanical fit.

- ① **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 36. Riser configurations: 1A + 2A**

Card type	Slot priority	Form factor
HWRAID BOSS	1	Low Profile
10Gb NICs	2, 3	Low Profile
HBA330 Mini	Integrated Slot	N/A

<b>Card type</b>	<b>Slot priority</b>	<b>Form factor</b>
rNDC	Integrated Slot	N/A

**Table 37. Riser configurations:1A**

<b>Card type</b>	<b>Slot priority</b>	<b>Form factor</b>
HWRAID BOSS	1	Low Profile
10Gb NICs (Intel)	2	Low Profile
HBA330 Mini	Integrated Slot	N/A
rNDC	Integrated Slot	N/A

# Getting help

Topics:

- [Contacting Dell EMC](#)
- [Registering for online support](#)
- [Accessing support resources](#)

## Contacting Dell EMC

You can link your Online Support account with **VxRail Manager** and access support resources without having to log in separately.

**NOTE:** If you plan to set up EMC Secure Remote Services (ESRS), you must link your Online Support account to VxRail Manager under the same ID or it may not work properly.

## Registering for online support

You can create an **Online Support** account to access support resources such as:

- Register your appliance.
- Obtain product license files and software updates.
- Download Dell EMC VxRail Series product documentation.
- Download the SolVe Desktop Procedure Generator.
- Browse the Dell EMC VxRail Series community and support information.
- Link your support account for access to resources from within **VxRail Manager**.

To register for online support:

- 1 Go to **[emc.com/vxrailsupport](https://emc.com/vxrailsupport)** or **[support.emc.com](https://support.emc.com)**.
- 2 Click **Register here**.
- 3 Fill in the required information.

You will receive a confirmation email within 48 hours.

## Accessing support resources

You can access support resources for your VxRail Series using one of the following methods:

- VxRail Manager Support
- **[emc.com/vxrailsupport](https://emc.com/vxrailsupport)** (or **[support.emc.com](https://support.emc.com)**)
- **<https://solve.emc.com>**

**NOTE:** Additional VxRail Series information is available through the SolVe desktop application. SolVe includes step-by-step procedures for replacing certain hardware components, and other tasks.

Entrar em contato com o suporte

VxRail Appliance Series

Add to My Saved Products Inscrever-se para atualizações do produto

Search Support or Find Service Request by Number

Advanced



O novo VxRail™ é um dispositivo de infraestrutura hiperconvergente que fornece virtualização, computação, armazenamento, sistema de rede e proteção de dados em um componente modular SDDC 2U com 4 nós e que é dimensionável e fácil de gerenciar. Os dispositivos VxRail são totalmente integrados e previamente configurados e testados com dispositivos VMware Mais...

- Visão geral do produto » Documentação » Conselhos » Downloads » Base de Conhecimentos » Ferramentas »

Recursos recomendados

- Top Service Topics - VxRail Key Documentation - VxRail

CONSELHOS

You are not receiving Advisory alerts for this product. Manage Advisory Subscription >>

Conselhos de segurança: 534941 Dell EMC VxRail Appliance 2019 Security Update for Multiple Intel MDS Vulnerabilities 21 de Junho de 2019 às 09:26

Conselhos de segurança: 533455 VxRail Security Update for Multiple Hardware Appliance Firmware Vulnerabilities 14 de Maio de 2019 às 09:49

Exibir tudo

SERVICE LIFE

Datas de versão e fim da vida útil

Series

- VxRail E Series Nodes VxRail G Series Nodes VxRail P Series Nodes VxRail S Series Nodes VxRail V Series Nodes

Products

- Pivotal Ready Architecture VxRail 460 and 470 Nodes VxRail Gen2 Hardware VxRail Software

DOCUMENTAÇÃO (70)

- VxRail Support Matrix 18 de Setembro de 2019 | Produto:VxRail Gen2 Hardware,VxRail Appliance Series,VxRail Software 4.0,VxRail G Series Nodes,VxRail G410...More | Tarefa de suporte:Install, Plan | 1,42 MB | pdf | Documento de compatibilidade VxRail Appliance Software 4.7.x Release Notes 18 de Setembro de 2019 | Produto:VxRail Appliance Series,VxRail G Series Nodes,VxRail G410,VxRail E Series Nodes,VxRail E460...More | Tarefa de suporte:Install, Plan | 0,17 MB | pdf | Notas da versão VxRail Appliance 4.5.x and 4.7.x API Guide 18 de Setembro de 2019 | Produto:VxRail Appliance Series,VxRail G Series Nodes,VxRail E Series Nodes,VxRail P Series Nodes,VxRail V Series Nodes...More | Tarefa de suporte:Administer, Using | 0,43 MB | pdf | Referência

More

BASE DE CONHECIMENTOS (702)

DOWNLOADS (58)

DISCUSSIONS (10)

Central de atendimento

- Ingressar no bate-papo on-line Criar chamado Entrar em contato com o Atendimento ao cliente

Expandir

- Transformação da TI Transformação da força de trabalho Produtos Soluções Serviços Histórias de clientes

- MyService360 Suporte por produto Downloads

Comunidade  
Centro de Serviços

---

Nossa liderança  
Sala de imprensa  
Parceria  
Legacy of Good  
Capital da Dell Technologies  
Investidores

---

Carreiras  
Fale conosco  
Programa de Parceiros Dell EMC

---

© 2019 Dell Inc. ou suas subsidiárias. Todos os direitos reservados.

[Privacidade](#) | [Termos de uso](#) | [Jurídico](#) | [Intel](#) | [Contato](#) | [Brasil](#)

# DELL EMC VXRAIL™ SYSTEM TECHBOOK

A Hyperconverged Infrastructure System from  
Dell EMC® and VMware®

## ABSTRACT

This TechBook is a conceptual and architectural review of the Dell EMC VxRail system, powered by VMware vSAN, and with Intel Inside®. The TechBook first describes how hyperconverged infrastructure drives digital transformation and then focuses on the VxRail system as a leading hyperconverged technology solution.

August 2019



# Contents

<b>IT's new challenge: Leading digital transformation</b>	<b>6</b>
Dell EMC digital transformation: Faster outcomes. Simpler IT	7
Accelerating IT transformation with Dell EMC hyperconverged infrastructure	7
Fully transforming to the hybrid cloud	8
Innovate rather than integrate	8
<b>Hyperconverged infrastructure: A modern infrastructure for modern IT challenges</b>	<b>10</b>
Enabling technologies for HCI	11
Drivers for hyperconverged infrastructure adoption	11
<b>Dell EMC hyperconverged infrastructure platforms</b>	<b>13</b>
Dell EMC HCI delivers a turnkey customer experience	14
<b>Dell EMC VxRail Systems</b>	<b>15</b>
VxRail Security and Compliance	17
<b>VxRail hardware architecture</b>	<b>19</b>
VxRail cluster	20
VxRail models and specifications (based on 14 <sup>th</sup> generation Dell EMC PowerEdge Servers)	20
VxRail node	21
Intel® Xeon® Scalable processor: Powerful processing for VxRail	23
VxRail node storage disk drives	23
VxRail hardware options	24
VxRail scaling	25
Upgradeable options	26
VxRail networking	26
1GbE network option	28
Dell EMC Open Networking & VxRail	29
Dell EMC SmartFabric Services (SFS)	29
<b>VxRail software architecture</b>	<b>31</b>
System management	31
VxRail HCI System Software	31
VxRail Manager	34
Customer upgradeable software	36
VxRail Analytical Consulting Engine (ACE)	37
vSphere and vSAN ordering information	38
<b>VMware vSphere</b>	<b>40</b>
VMware vCenter Server	40
vCenter services and interfaces	43
Enhanced Linked Mode	44
VMware vSphere ESXi	44
Communication between vCenter Server and ESXi hosts	45
Virtual machines	46
Virtual machine hardware	47
Virtual Machine Communication	47

Virtual networking .....	48
Virtual Distributed Switch.....	48
vMotion and Virtual Machine mobility .....	49
Enhanced vMotion Compatibility .....	51
Storage vMotion .....	51
vSphere Distributed Resource Scheduler .....	52
vSphere High Availability (HA) .....	54
vCenter Server Watchdog .....	56
vSphere Encryption .....	56
<b>vSAN .....</b>	<b>57</b>
Disk groups .....	58
Hybrid and All-Flash differences.....	59
Read cache: Basic function .....	59
Write cache: Basic function .....	59
Flash endurance .....	60
vSAN impact on flash endurance .....	60
Client cache .....	60
Objects and components .....	60
Witness .....	61
Replicas .....	61
Storage Policy Based Management (SPBM).....	62
Dynamic policy changes.....	62
Storage policy attributes .....	63
Sparse Swap .....	65
I/O paths and caching algorithms .....	66
Read caching.....	66
Anatomy of a hybrid read .....	67
Anatomy of an All-Flash read .....	68
Write caching .....	68
Anatomy of a write I/O—hybrid and All-Flash (FTM=mirroring).....	69
Distributed caching considerations.....	70
vSAN high availability and fault domains.....	71
Fault domain overview.....	71
Fault domains and rack-level failures .....	72
Cautions when deploying a minimum cluster configuration .....	73
vSAN Stretched Cluster.....	73
Stretched Cluster with Local Protection .....	74
Site locality.....	75
Networking.....	75
Stretched cluster heartbeats and site bias .....	75
vSphere HA settings for stretched cluster .....	75
2-Node Configuration.....	76
Snapshots .....	76
Storage efficiency using deduplication and compression.....	77
Deduplication and compression overhead .....	79
Erasure coding.....	80
Enabling Erasure Coding.....	81
Erasure coding overhead .....	81

vSAN Encryption.....	82
<b>VxRail integrated software .....</b>	<b>83</b>
VM Replication.....	83
VMware vSphere Replication .....	83
Dell EMC RecoverPoint for Virtual Machines .....	84
VxRail replication use case.....	86
<b>Support for external network storage .....</b>	<b>87</b>
Fibre Channel with VxRail .....	87
iSCSI with VxRail.....	87
NFS with VxRail .....	88
<b>VxRail solutions and ecosystem .....</b>	<b>90</b>
VMware Validated Design with VxRail .....	90
VMware Cloud Foundation on VxRail.....	92
Pivotal Ready Architecture (PRA) .....	93
Flexible consumption options .....	94
Flex on Demand - a cloud-like consumption option .....	94
VDI Complete .....	95
VMware Horizon .....	95
VMware Horizon with VxRail .....	96
VMware vSphere Platinum .....	97
IsilonSD Edge .....	97
SAP HANA Certification with VxRail.....	98
Reference Architecture for Splunk.....	99
<b>Additional Product information .....</b>	<b>101</b>
Dell EMC ProSupport for Enterprise.....	101
Dell EMC ProDeploy Services for VxRail Systems .....	101
Learn more.....	101
Contact.....	101
View more .....	101

## Preface

The Dell EMC TechBook is a conceptual and architectural review of the Dell EMC VxRail™ system, powered by VMware vSAN with Intel Inside. The TechBook describes how hyperconverged infrastructure drives digital transformation and focuses on the VxRail system as a leading hyperconverged technology solution.

## Audience

This TechBook is intended for Dell EMC field personnel, partners, and customers involved in designing, acquiring, managing, or operating a VxRail system solution.

## Related resources and documentation

Refer to the following items for related, supplemental documentation, technical papers, and websites.

Dell EMC VxRail Network Guide: <https://vxrail.is/networkplanning>

VxRail Planning Guide for Virtual SAN Stretched Cluster:  
<https://vxrail.is/stretchedclusterplanning>

An overview of VMware vSAN Caching Algorithms at  
<https://www.vmware.com/files/pdf/products/vsan/vmware-virtual-san-caching-whitepaper.pdf>

VMware vSAN 6.2 Space Efficient Technologies Technical White Paper at  
<http://www.vmware.com/files/pdf/products/vsan/vmware-vsan-62-space-efficiency-technologies.pdf>

VxRail Stretched Cluster at: <https://www.dell.com/resources/en-us/asset/offering-overview-documents/products/converged-infrastructure/vxrail-stretch-cluster-so.pdf>

vSphere Virtual Machine Administration Guide at  
<https://www.vmware.com/support/pubs/vsphere-esxi-vcenter-server-6-pubs.html>

vSphere Resource Management at <http://pubs.vmware.com/vsphere-65/topic/com.vmware.ICbase/PDF/vsphere-esxi-vcenter-server-65-resource-management-guide.pdf>

Dell EMC Hyperconverged infrastructure at <http://www.dell.com/en-us/work/shop/category/hyper-converged-infrastructure>

VxRail vCenter Server Planning Guide: <https://vxrail.is/vcenterplanning>

## IT's new challenge: Leading digital transformation

In the digital economy, applications are both the face and the backbone of the modern enterprise.

For the digital customer, user experience trumps all. Customer-facing applications must be available anytime, anywhere and on any device, and must provide real-time updates and intelligent interactions. For the business, the insights gleaned from the data collected from these interactions inform and drive future development needs.

Applications and the underlying infrastructure are strategic to the business. Businesses that can efficiently leverage modern datacenter technologies to rapidly deliver innovative capabilities to customers are positioned for real success.

The importance of applications in the modern enterprise presents a huge opportunity for IT organizations. No longer simply a back-office function, IT can lead a digital transformation that positions the business for success moving forward. IT can become an active enabler of the business.

Traditional IT teams are faced with a massive amount of complexity when building, configuring, maintaining and scaling applications. Organizations need to successfully deploy and operate an environment that takes full advantage of the innovation taking place across the industry – without the complexity of piecing together and supporting a wide range of patchwork tools.

The challenge is how to go about this transformation. Dell EMC surveyed<sup>1</sup> over 1,000 executives across multiple industries about the state of their digital transformation efforts. Survey questions focused on:

Modern datacenter technology utilization, such as the use of All-Flash arrays, scale-out architectures, converged and hyper converged platforms, and software-defined solutions across networking and storage domains.

Automated IT processes, as measured by the progress the organization has made in terms of running IT more like a public cloud provider (enabling self-service infrastructure provisioning; rapid scalability; and usage-based tracking and chargeback).

Transformed business and IT relationships: enabled by consistent communication between IT and business stakeholders and continuous inspection of IT outcomes by line of business (LOB) leadership.

The findings show that progress has been, at best, mixed. Some companies have barely started their digital transformation. Many have taken a piecemeal approach. Only a small minority have almost completed their digital transformation. Why is it taking so long?

The bottom line is that IT transformation is difficult. It requires a great deal of planning, evaluation, re-organization and modernization of infrastructure technologies and applications. Multiple factors including costs, skill sets, governance, the drive to innovate and willingness to transform influence whether a business moves beyond the traditional three-tier datacenter structure.

---

1

How IT Transformation Maturity Drives IT Agility, Innovation, and Improved Business Outcomes, Enterprise Strategy Group, April 2017. <https://www.emc.com/collateral/analyst-reports/esg-dellemc-it-transformation-maturity-report.pdf>

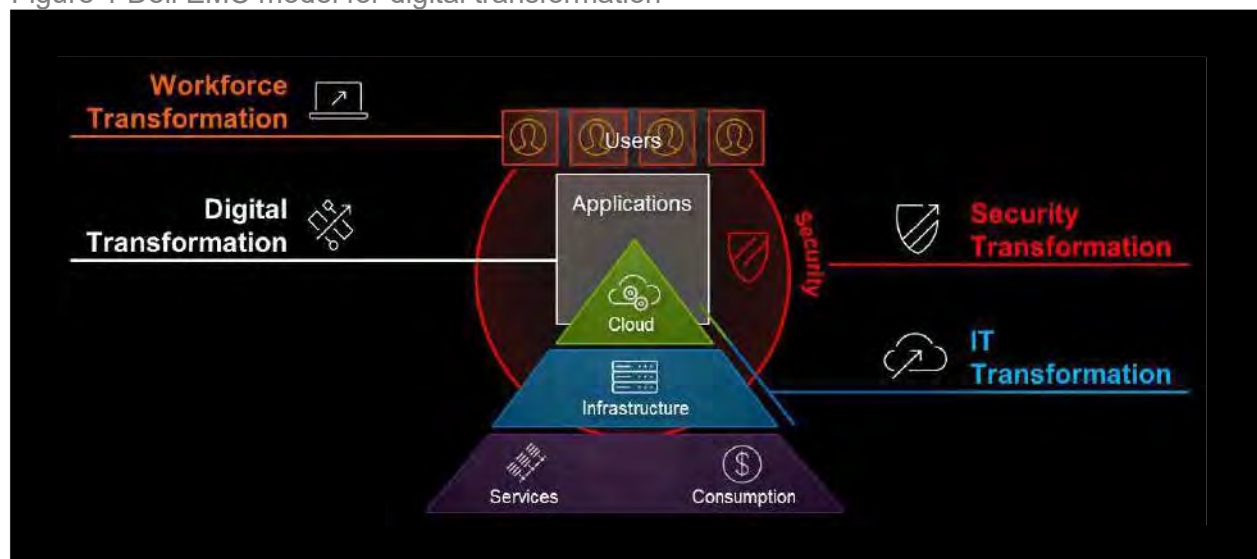
Every business approaches IT transformation at a different pace and has different goals for that transformation. Not every business wants or needs to go to a full cloud service delivery model. What is needed is an approach that enables businesses to transform to a place that provides the desired transformation benefits and at a pace that makes sense for their business model.

## Dell EMC digital transformation: Faster outcomes. Simpler IT.

To have a complete and strategic view of transformation, you must start with an application centric point of view. Ensuring that applications have the right strategy for service level agreements whether on premises or off premises entails having a hybrid cloud strategy. Dell EMC leverages hyperconverged infrastructure (HCI) to deliver completely integrated and supported HCI solutions and hybrid cloud platforms that modernize, automate, and transform the enterprise datacenter and IT Transformation.

The following figure shows the Dell EMC model for digital transformation.

Figure 1 Dell EMC model for digital transformation



Dell EMC delivers fully engineered turnkey hyperconverged infrastructure solutions that enable businesses to innovate faster and accelerate IT operations.

Dell EMC converged and hyperconverged infrastructure delivers application-focused solutions built on best-of-breed hardware and software that provide real business value, while dramatically reducing the risk and cost of deploying mission critical, general purpose, and cloud native applications.

And finally, Dell EMC offers a full range of services and flexible consumption models to help make it faster and easier to consume these solutions.

## Accelerating IT transformation with Dell EMC hyperconverged infrastructure

One of the first steps a business can take in their transformation journey is to simplify infrastructure deployment and management by introducing hyperconverged infrastructure (HCI) into the environment. HCI systems essentially collapse the traditional three-tier server, network, and storage model so that the infrastructure itself is much easier to manage.

Adopting hyperconverged infrastructure solutions that natively integrate compute, storage, virtualization, management and data services significantly reduces IT administrative tasks and create the foundation for a modern IT infrastructure. HCI solutions are optimal for reducing

infrastructure costs and simplifying management, regardless of workload deployment and extent of implementation.

### Fully transforming to the hybrid cloud

Many businesses would ultimately like to automate IT service delivery through a self-service catalog via a hybrid cloud. The hybrid cloud delivers the following benefits:

- A single control point for on- and off-premises resources
- Automation streamlines delivery of IT resources, delivering them in a consistent and repeatable manner aligned with business best practices
- Metering allows the IT team to communicate the value of services while providing visibility to the business on resource cost and consumption
- Self-service empowers application owners and business users to access the resources they need, when they need them
- Capacity management allows the IT team to better manage resources across the hybrid cloud
- Monitoring and reporting provides visibility to the capacity, performance and health of the environment
- Built-in security protects enterprise workloads
- Service-level choice aligns workloads to service levels and cost objectives
- Ability to meet the service level agreements with application level granularity

The vision of hybrid clouds is not new. Businesses have tried to deploy hybrid clouds using traditional infrastructure based on scale-up storage accessed over a storage network that is deployed and scaled in big chunks. While it is possible to build cloud capabilities on traditional three-tier infrastructure with scale-up storage, this is not the optimal solution.

If businesses want full IT transformation to the cloud support their application environment, Dell EMC can modernize, automate, and transform IT operations with complete turnkey, hybrid cloud platforms built on hyperconverged infrastructure.

### Innovate rather than integrate

Businesses do have the option of building a completely customized solution. Integrating storage, networking, compute, data protection, monitoring and reporting, and then figuring out how to get all of them to work together can be time consuming, but provides the most flexibility for an organization that may want prescribed vendor components as a part of their solution. Planning, designing and building a custom solution is a complex project that often takes months or years to come to fruition—too long if a business needs to roll out a solution to address immediate business needs and it can be costly to maintain or upgrade over the long term.

The challenge for IT is that complexity exists at each of these layers, so building and maintaining a functional, resilient cloud can be very difficult. Many companies find that doing it themselves requires more than 70% of their IT resources and budget, leaving few resources to focus on innovation and projects that add real value to the business.

For most businesses, the best way to consume HCI solutions is to buy them fully integrated with lifecycle management and single source of support. Buying versus building delivers the

best time-to-value, operational simplicity, and 5-year total cost of ownership savings of 619% over a traditional three-tier, build-your-own approach.<sup>2</sup>

---

2

Source: IDC Oct 2017

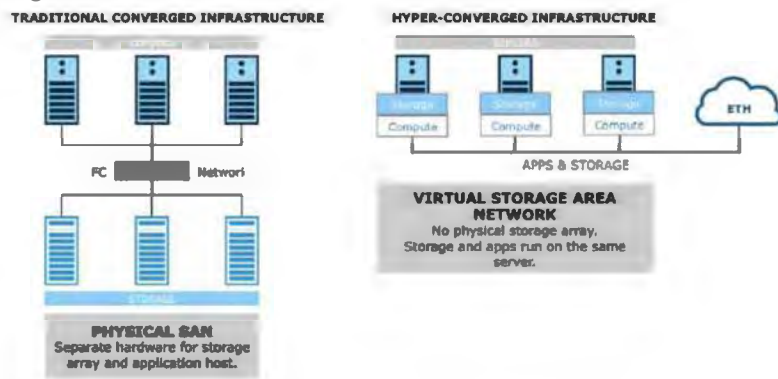


# Hyperconverged infrastructure: A modern infrastructure for modern IT challenges

*Converged infrastructure* platforms are fully pre-integrated server, traditional storage arrays, and networking hardware “stacks”. *Hyperconverged infrastructure* platforms are solutions that deliver compute, software-defined storage, and networking infrastructure services in a cluster of industry-standard servers.

Hyperconverged infrastructure extends the converged infrastructure model by incorporating the virtualization capabilities of software-defined storage (SDS). Hyperconverged infrastructure collapses the core components of traditional datacenter – compute and storage – into a server, effectively eliminating expensive and complex SAN environments. See the figures below.

Figure 2 CI and HCI



Because HCI is software-defined – which means the infrastructure operations are logically separated from the physical hardware – the integration between components is much tighter than with CI. HCI manages everything as a single system through a common toolset.

Hyperconverged infrastructure is particularly valuable because it lets you scale up quickly without a ton of added expense. That is not the case in traditional settings: customers either must buy more resources than they need in anticipation of scaling up, or wait until current workloads exhaust the allocated resources, then add infrastructure after the fact. Buying at the inopportune time means that resources are not optimally allocated and can even slow down your business from expanding.

HCI enables a pay-as-you-grow approach – start with what is needed today and expand incrementally rather than purchasing large amount of compute and storage up front. It also addresses the typical over-

## HCI deployment models

Dell EMC identifies four deployment models for IT transformation, defined by the business' desired end state and operational readiness to realize that end state.

### Build your own

Businesses wanting the benefits of HCI but prefer to maintain control & flexibility regarding server vendor and configuration choice.

### HCI systems

Systems are a more turnkey option, including the server, software stack, and lifecycle management. They are fully engineered and supported as a single product.

### Rack-scale HCI

Rack scale HCI solutions extend what is provided by HCI systems, and additionally, tightly integrates physical and software defined networking along with software defined storage, compute, virtualization, data protection, and management software to fully automate complete infrastructure configuration and provisioning through a software-defined management and orchestration layer. Like HCI systems they are engineered and supported as a single product.

### Turnkey HCI-based hybrid cloud

Combine either system or rack-scale HCI with additional cloud management, automation & orchestration software and pre-engineered services & workflows to deliver an agile, elastic, automated, self-sustaining private cloud that can integrate with off-premises public cloud providers.

provisioning and over-purchasing that occurs when technology is intended to last for multiple year cycles.

## Enabling technologies for HCI

The following table lists the confluence of technologies that has spurred the growth and development of hyperconverged infrastructure.

Table 1. Enabling technologies for HCI

Technology	Description
Software defined storage	<p>Abstracts the storage intelligence from the underlying storage infrastructure.</p> <p>Virtualizes direct-attach storage into a shared pool.</p> <p>Automates provisioning and load balancing.</p> <p>Allows a business to increase available storage resources, both capacity and processing power, by adding entire nodes (e.g., a server with storage software and media) to a cluster. The resulting cluster of nodes in turn acts as a single pool of storage capacity.</p>
Virtualization	<p>Abstracts compute and network functions.</p> <p>Enables physical resources to be shared.</p> <p>Improves utilization, mobility and security.</p>
X86 servers	<p>High performance processors, large memory.</p> <p>Flash media delivers consistent, predictable performance.</p>
Solid state storage	<p>Uses solid-state drives (most frequently various types of flash memory) to store data. This storage can reside in a storage controller or in a server, but for this assessment we are considering use cases limited to tiered and All-Flash storage arrays.</p> <p>In hybrid arrays, a portion of the drives in the array are solid-state and house the most active data on the array.</p> <p>In All-Flash arrays, all drives in the array are solid-state.</p>
High-speed networks	<p>Connects nodes together to create cluster.</p> <p>Enables HCI to deliver IOPS and reduced latencies.</p> <p>Connect applications to users</p>

## Drivers for hyperconverged infrastructure adoption

Customers that have transitioned or plan to transition to HCI state cost reduction, accelerated deployment, improved ability to scale, improved operational efficiencies and reduction in infrastructure tasks as top benefits they expect to realize when implementing HCI.


HCI delivers a compelling story for both CapEx and OpEx.

Savings in initial investments are lower, and operational expenses are also lower when compared to traditional three-tier architectures. Cost savings include power and cooling, ongoing system administration, and the elimination of disruptive upgrades and data migrations.

Rather than buying monolithic SAN-based infrastructure, a business can buy infrastructure that targeted for specific workloads. A main contributor to lower TCO and the increased agility of hyperconverged solutions is the ability start smaller and scale incrementally.

Not only is initial CapEx investment lower, but you can also scale them more incrementally, adding smaller amounts of compute or even expanding just the storage capacity as required, as shown in the figure below.


Figure 3 HCI: Buy what you need, pay as you go



Further, some HCI systems support multiple nodes in a single chassis. So, if space is available, you can simply plug a node into an existing chassis in a matter of minutes to scale your infrastructure. Deploying HCI reduces time, risk and complexity.

The following figure shows how HCI simplifies IT management and operations as compared to a traditional SAN environment.

Figure 4 HCI simplifies IT management and operations




The most compelling HCI solutions leverage management frameworks that drastically improve operational tasks, reduce the burdens of lifecycle management, and improve the level of responsiveness of IT staff to meet the needs of their business. Ideally, existing investments in management and orchestration stacks can be leveraged as well.

# Dell EMC hyperconverged infrastructure platforms

Dell EMC provides organizations the flexibility to choose HCI solutions that best fit their current state of IT transformation while ensuring IT certainty, continuous innovation and predictable evolution as they move toward cloud implementations. Whether a business is modernizing existing applications or deploying turnkey engineered solutions, the Dell EMC HCI portfolio delivers the power, simplicity, and certainty a business needs for the next phase of digital transformation.

The following figure shows the Dell EMC HCI portfolio, powered by Intel®.

Figure 5 Dell EMC HCI portfolio



The Dell EMC hyperconverged infrastructure portfolio includes both systems and rack-scale offerings and differentiates with both fully-integrated VMware-based solutions or opportunities to get turnkey outcomes through hypervisor choice, bare metal or multi-hypervisor options.

Systems accelerate the transformation of both the compute and storage layers for customers' datacenters by delivering turnkey outcomes on all-flash, software-defined, scale-out architectures. Rack-scale offerings offer additional transformation for those customers who are ready to fully modernize their datacenter by adopting software defined networking as well as compute and storage in a fully integrated, turnkey fashion.

Building upon the bedrock of the modern datacenter, Dell EMC delivers the #1 hyperconverged infrastructure portfolio purpose-built for HCI with the newest 14<sup>th</sup>-generation Dell EMC PowerEdge server platform. This portfolio delivers tailor-made performance and reliability powerful enough for any workload, combined with a polished approach to intelligent deployment and operations that simplifies and accelerates IT. Dell EMC HCI on latest generation PowerEdge servers are powerful, purposeful, and polished hyperconverged platforms that provide the ideal foundation for software-defined datacenter initiatives.

Customers demand more performance and reliability as HCI moves into the core datacenter where it must run a wider range of applications and workloads. With improved performance (double the IOPS), better economics (up to 3x more VDI users per node), and flexibility (more



configurability to meet more use cases), Dell EMC HCI on 14th Generation PowerEdge servers has enough power for any workload.

Dell EMC PowerEdge servers are purpose-built for HCI. With up to 150 customer HCI requirements built-in, PowerEdge servers are designed specifically for and tailored to HCI workloads that depend on the tenets of both servers and storage. This results in a more consistent, predictable, and reliable high-performing HCI that can meet any use case. With our comprehensive portfolio, Dell EMC can deliver the best fit for your customers' specific HCI needs – from workload requirements, to customer environment/standardization, to deployment preferences.

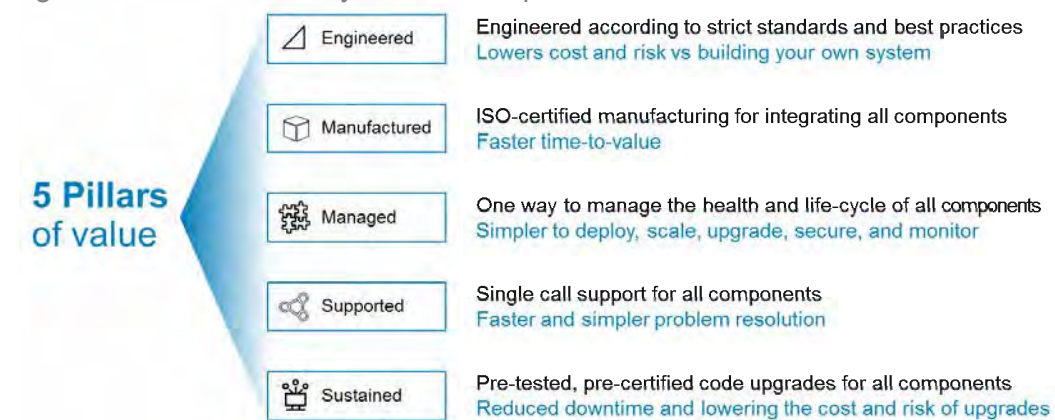
Dell EMC leads in hyperconverged sales with over 28% market share according to IDC<sup>3</sup>. More customers are choosing Dell EMC HCI over all others. Dell EMC PowerEdge is the world's best-selling server. Industry-leading Dell EMC HCI built on industry-leading PowerEdge, coupled with a single point of support and full lifecycle management for the entire system, makes for a winning solution.

## Dell EMC HCI delivers a turnkey customer experience

Dell EMC has invested thousands of engineering hours designing; integrating and testing its hyperconverged solutions to make sure all components are hardened, work together and are sustained as one. As a result, Dell EMC HCI solutions can be implemented in weeks rather than months.

The following figure identifies the key aspects of the Dell EMC HCI turnkey experience.

Figure 6 Dell EMC turnkey customer experience



Dell EMC is constantly innovating and delivering new capabilities to ensure its HCI solution continues to evolve to meet new customer requirements. New capabilities are added as part of the solution release cycles with Dell EMC testing the entire solution end-to-end, including upgrades from previous versions. In addition, major upgrades are supported by Dell EMC Global Services across the entire solution stack eliminating interoperability management concerns for IT Ops teams.

<sup>3</sup> Based on IDC Converged Tracker Q4 2018

# Dell EMC VxRail Systems

VxRail systems are jointly developed by Dell EMC and VMware and are the only fully integrated, preconfigured, and tested HCI system powered by VMware vSAN technology for software-defined storage. Managed through the ubiquitous VMware vCenter Server interface, VxRail provides a familiar vSphere experience that enables streamlined deployment and the ability to extend the use of existing IT tools and processes.

VxRail systems are fully loaded with integrated, mission-critical data services from Dell EMC and VMware including compression, deduplication, replication, and backup. VxRail delivers resiliency and centralized-management functionality enabling faster, better, and simpler

management of consolidated workloads, virtual desktops, business-critical applications, and remote-office infrastructure. As the exclusive hyperconverged infrastructure system from Dell EMC and VMware, VxRail is the easiest and fastest way to stand up a fully virtualized VMware environment.

VxRail is the only HCI system on the market that fully integrates Intel-based Dell EMC PowerEdge Servers with VMware vSphere, and vSAN. VxRail is jointly engineered with VMware and supported as a single product, delivered by Dell EMC. VxRail seamlessly integrates with existing (and optional) VMware eco-system and cloud management solutions, including vRealize, NSX, Horizon, Platinum and any solution that is a part of the vast and robust vSphere ecosystem.

VxRail provides an entry point to the software defined datacenter (SDDC) for most workloads. Customers of all sizes and types can benefit from VxRail, including small- and medium-sized environments, remote and branch offices (ROBO), and edge departments, as well as providing a solid infrastructure foundation for larger datacenters.

Small-shop IT personnel benefit from the simplicity of the system model to expedite the application-deployment process while still taking advantage of data services only typically available in high-end systems.

Larger datacenters benefit by rapid deployment where a complete vSphere environment can be installed and be ready to deploy applications within few hours of the system arriving on site. VxRail

allows businesses to start small and scale non-disruptively. Storage is configured to meet appropriate application capacity and performance requirements.

In addition, nodes are available with different compute power, memory, and cache configurations to closely match the requirements of new and expanding use cases. As



## VxRail essentials

Fully integrated, preconfigured, and tested hyperconverged infrastructure appliance simplifies and extends VMware environments

Seamlessly integrates with existing VMware eco-system management solutions for streamlined deployment and management in VMware environments.

Start small, with a few as three nodes. Single node scaling, storage capacity expansion, and vSphere license independence enable growth that meets business demands.

Backup distributed applications or workloads with integrated data protection options, including RecoverPoint for VMs.

Single point of global 24x7 support for both the hardware and software

requirements grow, the system easily scales out and scales up in granular increments. Finally, because the VxRail is jointly engineered, integrated, and tested, organizations can leverage a single source of support and remote services from Dell EMC.

VxRail environments are configured as a cluster consisting of a minimum of two server nodes, with each node containing internal storage drives. VxRail systems are delivered with the software loaded, ready to attach to a customer-provided network. While most environments use 10Gb Ethernet for internal and external communications, 25Gb or 1Gb Ethernet connectivity is also available. Using a simple wizard at the time of install, the system can be configured to match unique site and networking requirements.

VxRail systems enable organizations to start small and scale out as the IT organization transforms and adapts to managing converged infrastructure versus silos. With a rich set of data services, including data protection, tiering to the cloud, and active-active datacenter support, VxRail can be the foundational infrastructure for IT. Best of all, you can simply add new systems into existing clusters (and decommission aging systems) to provide an evergreen HCI environment, never having to worry about costly SAN data migrations ever again. As organizations continue to transform to a cloud model, integration with the VMware vRealize Suite enables full cloud automation and service delivery capabilities.

Dell EMC VxRail systems offer a choice of Dell EMC PowerEdge servers, powered by new Intel® Scalable® processors, variable RAM, and storage capacity, allowing customers to buy what they need now. The VxRail system uses a modular, distributed system architecture that starts with as few as two nodes and scales near linearly up to 64 nodes. Single-node scaling and storage capacity expansion provide a predictable, “pay-as-you-grow” approach for future scale up and out as business and user requirements evolve.

The VxRail software layers use VMware technology for server virtualization and software-defined storage. VxRail nodes are configured as ESXi hosts, and VMs and services communicate using the virtual switches for logical networking. VMware vSAN technology, implemented at the ESXi-kernel level, pools storage resources. This highly efficient SDS layer consumes minimal system resources, making more resources available to support user workloads. The kernel-level integration also dramatically reduces the complexities involved in infrastructure management. vSAN presents a familiar datastore to the nodes in the cluster and Storage Policy Based Management provides the flexibility to easily configure the appropriate level of service for each VM.

VxRail HCI System Software, the VxRail management platform, is a strategic advantage for VxRail and further reduces operational complexity. VxRail HCI System Software provides out-of-the-box automation and orchestration for day 0 to day 2 system-based operational tasks, which reduces the overall IT OpEx required to manage the stack. No build-it-yourself HCI solution provides this level of lifecycle management, automation, and operational simplicity. With VxRail HCI System Software, upgrades are simple and automated with a single-click. You can sit back and relax knowing you are going from one known good state to the next, inclusive of all the managed software and hardware component firmware. No longer do you need to verify hardware compatibility lists, run test and development scenarios, sequence and trial upgrades, and so on. The heavy lifting of sustaining and lifecycle management is already done for you.

Within the VxRail HCI System Software, the VxRail Manager plugin presents a simple integrated dashboard interface on vCenter Server for infrastructure monitoring and automation of lifecycle management tasks such as software upgrades. Since VxRail nodes function as ESXi hosts, vCenter Server is used for VM-related management, automation, monitoring, and security.

VxRail systems are powered by VMware vSAN software, which is fully integrated in the kernel of vSphere and provides full-featured and cost-effective software-defined storage. vSAN implements an efficient architecture, built directly into hypervisor. This distinguishes vSAN from solutions that typically install a virtual storage appliance (VSA) that runs as a guest VM on each host. Embedding vSAN into the ESXi kernel layer has advantages in performance and memory requirements. It has little impact on CPU utilization (less than 10 percent) and self-balances based on workload and resource availability. It presents storage as a familiar data store construct and works seamlessly with other vSphere features such as VMware vSphere vMotion.

vSphere is a well-established virtualization platform—a familiar usable entity in most datacenters. Dell EMC leverages vSphere for ESXi-based virtualization and VM networking in multiple product offerings, and they support a common set of VMware and Dell EMC services. This enables a VxRail implementation to integrate smoothly into VMware-centric datacenters and to operate in concert with Dell EMC converged, hyperconverged, and traditional storage offerings. NSX for SDN can optionally be added to the VxRail solution. NSX for SDN allows all datacenter assets to be maintained using a single administrative platform. Monitoring, upgrading, and diagnostics activities are performed efficiently and reliably.

Additional data services integrated into VxRail include RecoverPoint for VM replication and Dell EMC Remote Secure Services (ESRS).

To summarize, VxRail creates IT certainty. It is built on the enterprise-proven VMware hypervisor, powered by vSAN technology, delivers simplified lifecycle and daily operational management, and supports the dynamic and robust vSphere ecosystem. The system architecture is designed to predictably evolve with the business. The system scales by deploying preloaded systems to effortlessly extend the environment, rather than by building and manually integrating servers.

VxRail, powered by vSAN and featuring Intel-based Dell EMC PowerEdge Servers, continues to evolve so business can thrive without worrying about the IT infrastructure. VxRail is a fully integrated, preconfigured, and pre-tested VMware hyperconverged infrastructure system family. Based on VMware's vSphere and powered by vSAN and Dell EMC software, the VxRail system provides the IT infrastructure foundation for digital transformation that empowers organizations to continuously innovate.

## VxRail Security and Compliance

Dell EMC VxRail system is a resilient, secure, and modern hyperconverged infrastructure system that directly addresses the challenges of security and compliance in modern day environments.

VxRail system is engineered, built, configured, and maintained following the Dell EMC Secure Development Lifecycle, which follows a rigorous approach to secure product development, including executive-level risk management before products are shipped to market. Additionally, VMware vSphere—a significant part of VxRail hyperconverged infrastructure—has also been developed using a similar Security Development Lifecycle.

Everything that comprises VxRail is secure and can be seen in the figure below. Each component has security built in, with corporate security processes, unique security features, and supply chain control, so you can feel confident that VxRail can fit into your secure IT infrastructure design. The hardware is comprised of Dell EMC PowerEdge servers and Intel processors. The virtualization and software layers are comprised of vSphere and vSAN which is integrated into the kernel of vSphere. The integrated software and management included with



VxRail is comprised of software from VMware, the vRealize Log Insight and vCenter, and software from Dell EMC, RP4VM, ESRS, and the VxRail HCI System Software for all the lifecycle management of everything in this stack. (RP4VM excluded from LCM.) All of this is jointly engineered with Dell EMC and VMware, and delivered by and supported exclusively by Dell EMC as a single product—VxRail.

Figure 7 What comes with VxRail



VxRail is designed to a number of standards, has attained the Common Criteria EAL2+certificate, USGv6 certification making it IPv6 Ready, and provides a VxRail Product Security Configuration Guide to further harden VxRail deployments. Additionally, customers can leverage the VxRail STIG Compliance Guide and automated scripts to further harden their environments.

To learn more about VxRail’s Comprehensive Security by Design, please download the whitepaper: <https://www.emc.com/collateral/white-papers/vxrail-comprehensive-security-design.pdf>

## VxRail hardware architecture

The Dell EMC VxRail family is the standard in hyperconverged infrastructure, providing extreme flexibility to granularly add capacity and performance on demand and enabling customers to easily extend use cases across the VMware virtualized environment. The system-based design allows IT centers to scale capacity and performance non-disruptively, so they can start small and grow incrementally with minimal up-front planning. VxRail environments can be designed to support a small number of virtual machines and scale to thousands.

The VxRail architecture enables a predictable pay-as-you-grow approach that aligns to changing business goals and user demand. Dell EMC and VMware are continuously innovating, and VxRail introduced new Dell EMC PowerEdge-based models that offer extreme configuration flexibility. This flexibility allows customers to choose performance, graphics, and capacity as required for VMware environments, and supports more use cases.

The Dell EMC VxRail family of systems offers a range of platforms:

**G Series**—general-purpose and compute dense, multi-node form factor, ideal for widely deployed general purpose applications and VDI workloads (not requiring GPU cards)

**E Series**—everywhere from datacenter core to edge deployments, the combination of density, drive groups, and balance of resources in a low profile 1U form factor enables it to be deployed for a wide range of use cases.

**P Series**—performance optimized for high-end use cases with business-critical, performance-intensive applications and/or in-memory databases.

**V Series**—VDI optimized for specialized use cases with graphics intensive applications such as high-end 2D/3D visualization applications; the only series that supports GPU cards.

**S Series**—storage-dense configurations targeted at specialized use cases that require higher storage capacity at the server level such as Big Data, analytics, or collaboration applications.

The E, P, V, and S Series are single-node systems based on Dell EMC PowerEdge server technology, the number-one selling X86 server platform, with greater storage capacity, larger memory, and more powerful CPU options. The G Series a four-node system in a compact 2U chassis, providing a compute dense footprint.

VxRail systems are built using a distributed-cluster architecture consisting of modular blocks that scale linearly as the system grows from as small as three nodes to as large as 64 nodes. Nodes are available with different form factors, with single-node systems for use cases: low-profile systems; performance optimized; VDI optimized with GPU; and storage-optimized configurations supporting high-capacity HDD drives.

Extensive compute, memory, and storage options are designed to fit multiple use cases. Customers can choose from a range of next-gen Intel processors, variable memory sizes, storage, and cache capacity to provide the right balance of compute, memory, and storage. Single-node scaling and a low-cost entry point let customers buy just the right amount of storage and compute for today's requirements and effortlessly scale to accommodate tomorrow's growth. Systems are available with all-flash storage configurations that deliver the industry's most powerful HCI for applications that demand maximum performance and low latency.

## VxRail cluster

VxRail nodes are enclosed in a one-node, single server system, with each node having one or two multi-core processors and either all-flash solid state disks (SSDs) or a hybrid mix of flash SSDs and hard disk drives (HDDs). The nodes form a networked cluster with a minimum of three nodes and a maximum of 64. Nodes within a cluster must be of the same storage configuration, either all hybrid or all-flash. The flexibility to mix nodes within a cluster is supported. The first three nodes must have the same compute, memory, and storage configuration, and mixing 1GbE, 10GbE, and 25GbE is not supported. From the minimum configuration to the maximum, the VxRail cluster is easily expanded one node at-a-time.

System models support either 25GbE, 10GbE or 1GbE network. 10Gb and 25Gb Ethernet networks are required for all-flash configurations and environments that will scale to more than eight nodes. Additional ports are available, allowing the customer to expand VM-network traffic.

## VxRail models and specifications (based on 14<sup>th</sup> generation Dell EMC PowerEdge Servers)

VxRail systems built on the new 14th Generation Dell EMC PowerEdge server platform deliver the performance and reliability your customers need for the widest range of workloads, all with full lifecycle management from a single point of support. In short, VxRail is the fastest and easiest way to transform infrastructure. It takes a lot of work and expertise to engineer a high performance and reliable HCI solution, and the work does not stop after the initial deployment. Continuous validation is needed to keep it running smoothly through software upgrades and node additions. As a turnkey, pre-integrated, tested, and validated HCI solution, VxRail can be quickly deployed, easily distributed, and counted on to increase the predictability, availability, and performance of your IT environment.

VxRail systems on next generation servers include multiple purpose-built platforms with build-to-order configurations that support a wide range of customer use cases, including graphics-intensive VDI, big data and analytics, high performance computing, remote office, and more. With more processor options, new SATA SSDs, more additional network connectivity options, and more GPU expansion, you can now more closely match a VxRail to your workload requirements. No over provisioning here: buy what is needed, when it is needed.

VxRail models are available to meet the requirements of a wide set of use cases. For smaller workloads, there is a low-profile system space efficient configuration that uses 1U single-node systems. A performance-optimized and a VDI-optimized configuration is available in both all-flash and hybrid configurations. For use cases requiring even greater storage, a hybrid storage-dense configuration that uses larger-capacity 3.5-inch drives is available. All models have a wide range of available memory, SSD cache, capacity storage configuration options and can start as with as few as two nodes.

The VxRail on 14th generation PowerEdge servers is now more powerful and predictable than the previous generation with utmost flexibility to meet any use case and more demanding workloads in VMware environments.

Double the IOPS, half the response times – VxRail is powerful enough for anything






Sustained response time less than 1ms – VxRail delivers highly predictable performance

Dell EMC offers the world’s most configurable HCI systems – VxRail perfectly matches any HCI requirements



The following figure shows the range of platforms designed to support multiple use cases.

Figure 8 VxRail based on 14th generation Dell EMC PowerEdge servers


G Series Nodes	E Series Node	P Series Nodes	V Series Nodes	S Series Nodes
				
Compute dense	Low profile	Performance optimized	VDI optimized	Storage dense
G560/F/N	E560/F/N	P570/F/N	V570/F	S570
<b>Supporting latest generation to Dell EMC PowerEdge servers based on Intel® Xeon® Gen 1 and 2 Processor Scalable Family</b>				
Higher core counts, faster clock frequency, more memory channels, faster memory, higher endurance and redundant boot devices				
2000W or 2400W PSU 10GbE Optane & NVMe cache Mixed-use SAS cache	1100W PSU 10GbE or 25GbE Optane & NVMe cache Mixed-use SAS cache FC HBA	1100W or 1600W PSU 10GbE or 25GbE Optane & NVMe cache Mixed-use SAS cache FC HBA	2000W PSU 10GbE or 25GbE Up to 6 GPUs Mixed-use SAS cache FC HBA Dual Socket Only	1100W PSU 10GbE or 25GbE Hybrid Only Mixed-use SAS cache FC HBA

## VxRail node

The VxRail system is assembled with proven server-node hardware that has been integrated, tested, and validated as a complete solution by Dell EMC. All the nodes in the current generation of VxRail use Intel Xeon Scalable family processors. The Intel Xeon Scalable family processors feature a multi-threaded, multi-core CPU designed to handle diverse workloads for cloud services, high-performance computing, and networking. The number of cores and memory capacity differ for each VxRail model.

The figure below shows a physical view of a node server with its processors, memory and supporting components. All VxRail models have similar components but may be physically laid out differently.

Figure 9 VxRail P Series node server: back view



Each node server includes the following technology:

One or two Intel Xeon Scalable processors with up to 28 cores per processor

Up to 24 DDR4 DIMMs, providing memory capacity ranging from 64GB to 3072GB per node, depending on model

A PCIe SAS disk-drive controller supporting 12GB SAS speeds

A mirrored pair of BOSS 240GB SATA M.2 cards used to boot ESXi on the node

Four port 10GbE Network Daughter Card (auto-negotiable to 1GbE)

iDRAC port



## Intel® Xeon® Scalable processor: Powerful processing for VxRail

Intel® Xeon® Scalable platforms are powerful infrastructure that represents an evolutionary leap forward in agility and scalability. Disruptive by design, it sets a new benchmark in platform convergence and capabilities across compute, storage, memory, network and security. An innovative approach to platform design in Intel® Xeon® Scalable processors unlocks the power of scalable performance for today's datacenters and communications networks—from the smallest workloads to your most mission-critical applications.

With up to 28 cores delivering highly enhanced per core performance, and significant increases in memory bandwidth (six memory channels) and I/O bandwidth (48 PCIe lanes), your most data-hungry, latency-sensitive applications such as in-memory databases and high-performance computing will see notable improvements enabled by denser compute and faster access to large data volumes. And the latest generation processors designated with an 'M' can support denser memory, with up to 1536GB per processor.

The convergence of compute, memory, network, and storage performance combined with software ecosystem optimizations make Intel® Xeon® Scalable platforms ideal for fully virtualized, software-defined datacenters that dynamically self-provision resources—on-premises, through the network, and in the public cloud—based on workload needs.

### VxRail node storage disk drives

Storage capacity for the VxRail system is provided by disk drives that have been integrated, tested, and validated by Dell EMC. Most VxRail configurations use 2.5" form-factor SSDs and mechanical HDDs, and a configuration that uses 3.5" form-factor drives is also available for dense-storage requirements. Disk drives are logically organized into disk groups. Disk groups are configured in two ways:



### Intel Inside®. Trusted clouds outside.

Intel innovation is driving the modernization and hybrid cloud transformation of the traditional enterprise datacenter.

Migrating to the newest generation of high-performing and energy-efficient Intel-based hardware tunes a datacenter for highly optimized performance across a broad set of enterprise workloads while lowering costs and improving resource utilization.

Over time, evolving to a software-defined infrastructure (SDI) across all the critical domains of the datacenter (compute/storage/network) will deliver critical automation, orchestration and telemetry capabilities to help businesses unlock the full capabilities of multi-cloud computing.

With modern, industry-standard Intel® servers and technologies that run on software-defined infrastructure, you can seamlessly manage an environment that supports development and delivery of cloud-native applications and mission-critical workloads on secure private clouds, while also integrating with public clouds, many of which already run on Intel® architecture.

Hybrid configurations, which contain a single SSD flash-based disk for caching (the cache tier) and multiple HDD disks for capacity (the capacity tier)

All-flash configurations, which contain all SAS SSD or NVMe drives for both cache and SAS or SATA SSD for capacity. NVMe is supported in dual processor configurations only.

The flash drives used for caching and capacity have different endurance levels. Endurance level refers to the number of times that an entire flash disk can be written every day for a five-year period before it has to be replaced. A higher-endurance SSD is used for write caching, and capacity-optimized SSDs are used for capacity. Currently, the 400GB, 800GB, and 1600GB SAS SSD or 1600 GB NVMe are used for caching, and for capacity either 1.92 or 3.84TB flash SSDs, 1.2, 1.8, and 2.4 TB 10K HDDs, 2TB (7.2K) HDDs, and 4TB 7.2K (3.5" form-factor) are used. All VxRail disk configurations use a carefully designed cache-to-capacity ratio to ensure consistent performance. Capacity SSDs are offered in both higher endurance SAS and SATA. The SATA SSDs are a lower cost option, up to 30% per drive, and great for read or moderately intensive workloads. The figure below depicts the wide-ranging set of components available across the VxRail system line-up.


### VxRail hardware options

VxRail nodes can be configured with choice of processor, memory, storage (cache and capacity drives), based networking, power supply, and GPU (in the V Series only). The figure below shows the comprehensive set of options available across the family. Customers can be assured their VxRail is configured to best match their workload requirements in a very prescriptive manner, with millions of possible configuration combinations in the VxRail Series. Combine this with the numerous ways to scale on demand, and it is clear that VxRail provides the agility demanded by today's modern IT. Some of the upgrade options for VxRail include, memory, GPU, NIC cards, cache SSD, and capacity drives.

Figure 10 A representative set of component options available across the VxRail series

## VxRail Configuration Flexibility for Your Workload


G, E, P, S, V Series based on the latest Dell EMC PowerEdge servers



## VxRail scaling

VxRail clusters start with as few as three nodes and can grow in one-node increments up to 64 nodes, providing performance and capacity to meet a wide range of use cases. Two node VxRail clusters are supported but cannot be expanded at this time. New systems can be added non-disruptively and different models can be mixed within a VxRail cluster. Flexible storage options also allow a node to start with a few drives and add drives as capacity requirements grow, as shown in the following figure. Single node upgrades and drive scalability protect an optimized initial investment by allowing customers to start with what they need and expand the VxRail cluster by adding nodes and/or drives to increase performance and capacity as needed. Consult your Dell EMC representative for assistance.

Figure 11 VxRail scale on demand



A few basic rules regarding scaling are worth considering for planning:

### 1. Balance

- The first three nodes in a cluster must be the same processor and drive configuration. (2-Node vSAN configurations are not supported.)
- All nodes must be running the same version of software.
- Cannot mix hybrid and all flash nodes in the same cluster.
- 1GbE, 10GbE, and 25GbE base networking cannot be mixed in the same cluster.
- 1GbE must be hybrid and single processor node type.
- For G Series, all nodes in a chassis must be identical.

### 2. Flexibility

- Systems in a cluster can be different models or series and can have different numbers of nodes.
- A cluster can have a varied number of drives, CPU, memory, and model types.
- A cluster can have between 3-64, but only a max of 8 if 1GbE networking is used.
- For G Series, a chassis can be partially populated.



## Upgradeable options

With VxRail, nodes can upgrade or add memory, NIC cards, cache drives, and capacity drives. GPU can be upgraded or added in the V Series. It is not possible to upgrade from a single processor to a dual processor VxRail node. Please refer to the following table for information on which components are customer installable (replaceable).

Figure 12 VxRail customer and field replaceable parts

Hardware Component	Customer Replaceable Unit (CRU)	Field Replaceable Unit (FRU)
System Memory	Y	
Hard Drive	Y	
Solid State Drive (cache and capacity)	Y	
NVMe Cache Drive		Y
PCIe Network Interface Cards	Y	
Graphical Processing Unit (GPU)	Y	
Micro SDHC Card	Y	
Power Supply	Y	
Processors		Y
System Motherboard		Y
Host Bus Adapter (HBA330)	Y	
BOSS controller card and M.2 SATA disk		Y
Network Daughter Card (NDC)		Y

<sup>a</sup>The above table is a non-exhaustive list of FRUs that reflects common top level assembly parts.

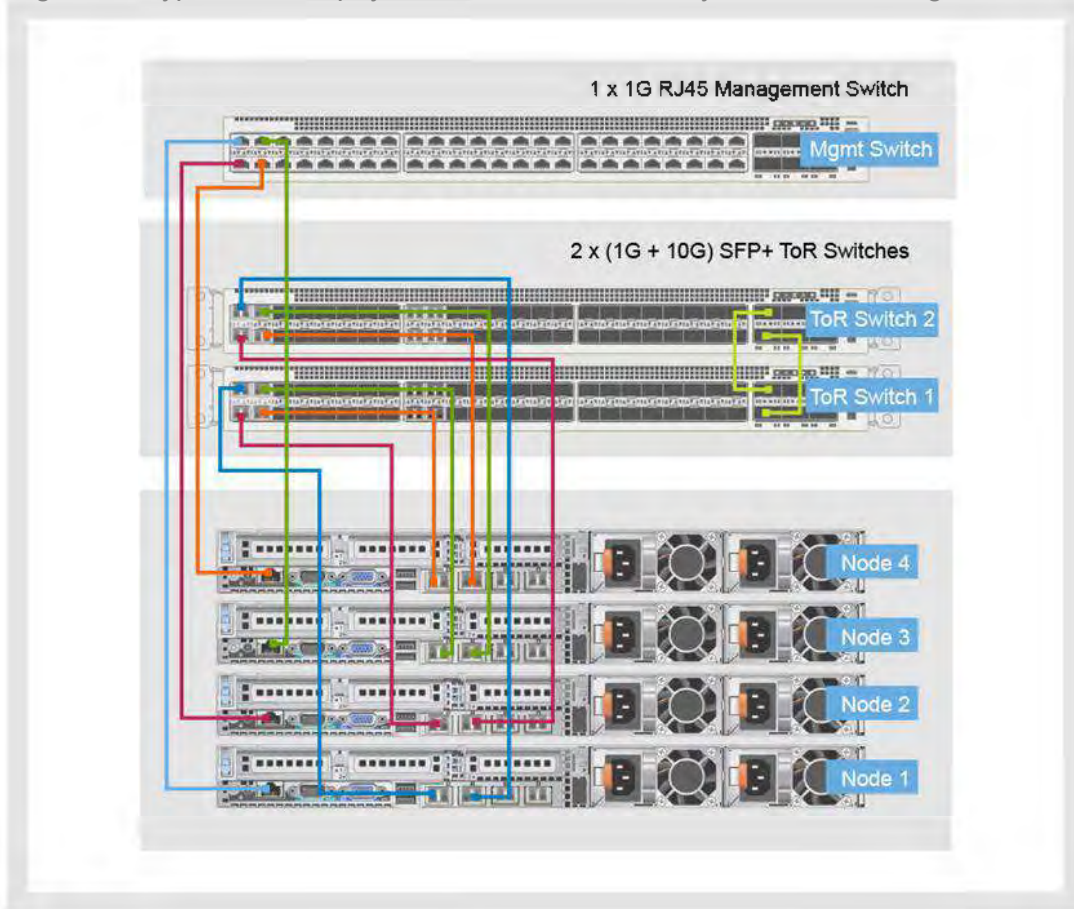
## VxRail networking

The VxRail system is a self-contained environment with compute, storage, server virtualization, and management services that make up a hyperconverged infrastructure. The distributed cluster architecture allows independent nodes to work together as a single system. Each node both contributes to and consumes system resources. This close coupling between nodes is accomplished through IP networking connectivity. IP networking also provides access to virtual machines and the services they provide.

While VxRail is a self-contained infrastructure; it is not a stand-alone environment. It is intended to connect and integrate with the customer's existing datacenter network. A typical implementation uses one or more customer-provided 10GbE Top of Rack (ToR) switches to connect each node in the VxRail cluster. For smaller environments, an option to use 1GbE switches is available, but these lower-bandwidth networks limit performance and scale. While the network switches are typically customer provided, Dell EMC offers an Ethernet switch, S4048, which can be included with the system.

The figure below shows typical network connectivity using two switches for redundancy. Single-switch implementations are also supported.

Figure 13 Typical VxRail physical network connectivity for 10GbE configurations



The number of Ethernet switch ports required depends on the VxRail model and whether it is configured for hybrid storage or for all flash. The all-flash system requires two or four 10GbE ports, and hybrid systems use either two 10GbE ports per node or four 1GbE ports per node. For 1GbE networks, the 10GbE ports auto-negotiate down to 1GbE. A two port 25GbE SFP28 is also an option. Additional network connectivity can be accomplished by adding additional NIC cards. The additional PCIe NICs are not configured by VxRail management but can be used by the customer to support non-VxRail traffic, primarily VM traffic. The additional ports are managed through vCenter.

Network traffic is segregated using switch-based VLAN technology and vSphere Network I/O Control (NIOC). Four types of network traffic exist in a VxRail cluster:

**Management.** Management traffic is used for connecting to VMware vCenter web client, VxRail Manager, and other management interfaces and for communications between the management components and the ESXi nodes in the cluster. Either the default VLAN or a specific management VLAN is used for management traffic.

**vSAN.** Data access for read and write activity as well as for optimization and data rebuild is performed over the vSAN network. Low network latency is critical for this traffic and a specific VLAN isolates this traffic.

**vMotion.** VMware vMotion™ allows virtual-machine mobility between nodes. A separate VLAN is used to isolate this traffic.

**Virtual Machine.** Users access virtual machines and the service provided over the VM network(s). At least one VM VLAN is configured when the system is initially configured, and others may be defined as required.

Pre-installation planning includes verifying that enough physical switch ports are available and that the ports are configured for the appropriate VLANs. VLANs along with IP addresses and other network-configuration information are used when the system is configured during installation. Detailed planning and configuration information is included in the [VxRail Network Guide](#).

When the system is initialized during installation, the configuration wizard automatically configures the required uplinks following VxRail standards and best practices. The wizard asks for the NIC configuration and accepts two options:

**4X1GbE.** Only valid for systems with hybrid-storage configuration with a single processor. The four (4) 10GbE ports auto-negotiate down to 1GbE. Management, vSAN, vMotion, and VM traffic is associated with these ports with the appropriate network teaming policy and NIOC settings.

**2X10GbE.** Management, vSAN, vMotion, and VM traffic is associated with these ports with the appropriate network teaming policy and NIOC settings. Note the additional two ports are available to use for additional VM traffic.

**4X10GbE.** Management, vSAN, vMotion, and VM traffic is associated with these ports with the appropriate network teaming policy and NIOC settings.

**2X25GbE.** Management, vSAN, vMotion, and VM traffic is associated with these ports with the appropriate network teaming policy and NIOC settings.

If nodes have additional physical NIC ports, they can be configured after installation using standard vSphere procedures.

### 1GbE network option

The Ethernet network not only provides connectivity to the VMs and services, it also provides the backplane for the nodes in a hyperconverged infrastructure to aggregate and share system resources. Therefore, network bandwidth is critical to system scale and performance. Today, most datacenters are built with 10Gb Ethernet connectivity, but 1GbE still exists in many environments. To support these environments, Dell EMC also offers 1GbE VxRail models for smaller, less-demanding workloads. The following are considerations for using the 1GbE connectivity option:

1GbE is only supported for hybrid-storage configurations, as it does not provide the bandwidth necessary for the ultra-high performance required by an all-flash system.

The maximum supported node count is eight nodes per cluster, because vSAN traffic increases with the number of nodes.

Only nodes with single-socket CPUs are supported.

A minimum of four ports are required per node. This increases the total number of switch ports required.

## Dell EMC Open Networking & VxRail

As hyperconverged clusters scale, the network fabric becomes the critical piece of a successful deployment. Dell EMC fabrics deliver:


- On-demand scalability: to remain competitive, the modern datacenter requires the capability to dynamically grow and shrink based on business requirements. Together, Dell EMC Networking switching fabrics and Dell EMC VxRail provide an intelligent and capable architecture that scales on demand and increases the efficiency of the datacenter.
- Increased availability at scale: robust and redundant fabrics and storage are an absolute necessity for today's datacenter. A single failure should not cause a full-service interruption.

Dell EMC Fabric solutions provide a proven, fully-integrated and adaptable infrastructure. Our solutions leverage components of a hyperconverged architecture and purpose-built networking infrastructure elements, created solely to address the requirements of hyper-convergence and efficient data consumption.

### Dell EMC SmartFabric Services (SFS)

Dell EMC Networking and VxRail are driving products and solutions by offering a well-engineered solution for the software defined enterprise that can deliver both operational and infrastructure efficiencies that were previously unavailable. Dell EMC is the only vendor with a complete set of hardware and software products capable of providing the necessary tools for the digital transformation. With the introduction of VxRail 4.7 and Dell EMC Networking OS10 Enterprise Edition SmartFabric Services, the conversation around automated dynamic infrastructure deployment can finally take place. Dell EMC, with VxRail and SmartFabric, is the first and only solution to deliver fully automated network awareness and configuration during set up, cluster expansion, and during day-to-day management to help create IT certainty.

Figure 14 SmartFabric Services



SmartFabric Services is a “one-of-a-kind” feature, part of the Dell EMC OS10 Enterprise Edition flagship networking operating system. Its introduction creates a fully integrated solution between the fabric and a hyperconverged cluster infrastructure such as VxRail. With SmartFabric Services, customers can quickly and easily deploy and automate datacenter networking fabrics. This enables faster time to production for hyperconverged and private cloud environments at any scale while being fully interoperable with existing datacenter infrastructure.

## VxRail software architecture

The prior section introduced the flexible VxRail hardware architecture based on proven server technology. While the VxRail hardware helps differentiate the VxRail from other HCI solutions, VxRail is a complete system that includes software that enables a software-defined datacenter. These sections on software architecture provide a comprehensive examination of all the VxRail software components and their relationships and co-dependencies.

The VxRail system is architected with software stack for system management, virtualization, and VM management. The stack comes pre-installed and simply requires running a configuration wizard on site to integrate the system into an existing network environment. VxRail HCI System Software is included for system management, operations, and automation. The VMware virtualization and virtual infrastructure management software includes:

- VMware vCenter Server
- vSphere ESXi
- vSAN (software-defined storage)
- VMware vRealize™ Log Insight™

Additional Dell EMC software includes:

RecoverPoint for VMs—5 VM licenses per node (for single node systems), 15 VM per chassis for the G Series

VxRail provides a unique and tightly integrated architecture for VMware environments. VxRail deeply integrates VMware virtualization software. Specifically, VMware vSAN is integrated at the kernel level and is managed with VMware vSphere, which enables higher performance for the VxRail system as well as automated scaling and wizard-based upgrades.

### System management

The introduction section of this TechBook addresses the complexity of the software-defined datacenter and the challenges of managing and maintaining an SDDC environment, and VxRail Manager responds to this challenge directly.

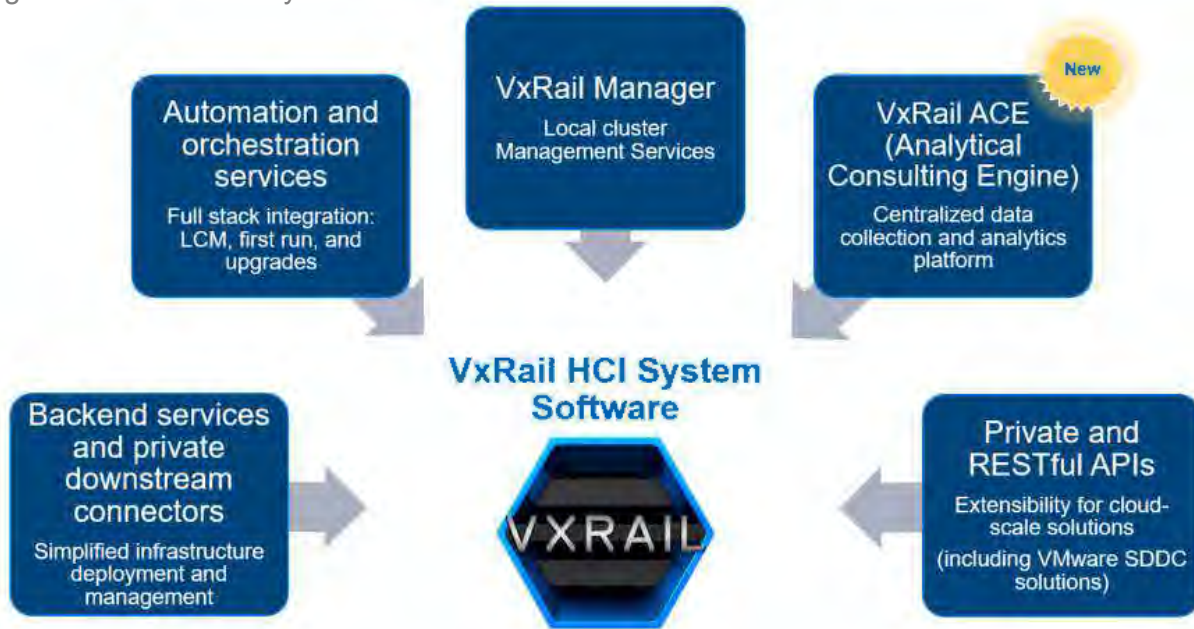
#### VxRail HCI System Software

VxRail HCI System Software, the VxRail management platform, is the system hardware lifecycle management and serviceability interface for VxRail clusters. It is a strategic advantage for VxRail and further reduces operational complexity. VxRail Manager provides out-of-the-box automation and orchestration for day 0 to day 2 system-based operational tasks, which reduces the overall IT OpEx required to manage the stack. No build-it-yourself HCI solution provides this level of lifecycle management, automation, and operational simplicity.

All virtualization management is performed using vCenter.  
VxRail HCI System Software does not do any virtualization management.



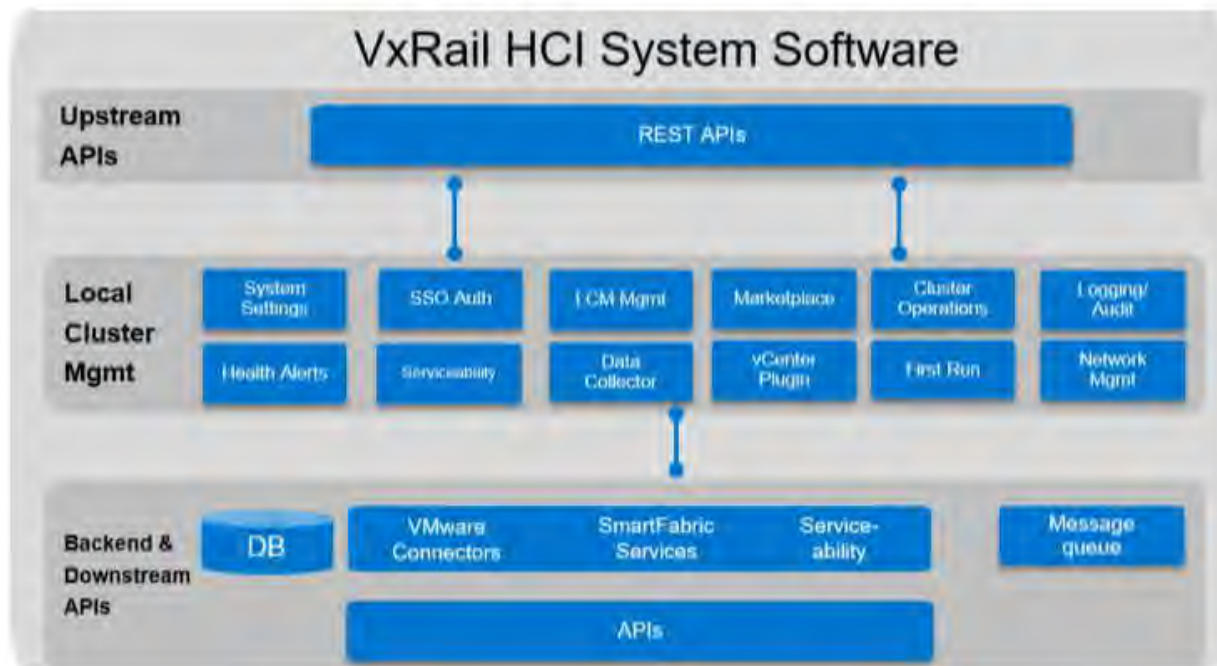
Figure 15 VxRail HCI System Software



To provide this level of differentiated value to customers, VxRail HCI System Software relies on a set of backend services to gather information and provide event coordination with the underlying infrastructure components including the vSAN cluster (i.e. ESXi, vSAN, vCenter Server), Dell PowerEdge server, and serviceability platforms (i.e. ESRS and eServices). This integration allows VxRail to automate and orchestrate infrastructure processes into critical services highly valued by HCI customers (i.e. lifecycle management, single package hardware/software upgrades, automated deployment capabilities). Users can access these cluster management services from their vCenter Server with the VxRail Manager plugin, or via a set of RESTful APIs. By providing public APIs, the value of VxRail for HCI can extend to SDDC solutions such as VMware Cloud Foundation, custom cloud solutions (i.e. Puppet, Ansible), or scripted solutions if customers, like service providers, wish to deploy and manage VxRail clusters at scale.

New to the VxRail management platform is another critical factor to customers' journey in IT transformation. While operational simplicity heavily impacts OpEx, the rise in infrastructure machine learning is fast becoming an essential component to aid in IT transformation. VxRail Analytical Consulting Engine (ACE) is the newly introduced analytical platform that leverages the data collection from VxRail clusters and best practices to deliver operational intelligence to customers about their HCI environment.

Figure 16 Architecture of VxRail HCI System Software



The architecture diagram above expounds on VxRail HCI System Software. From the bottom up, the backend services include APIs that connect downstream to various members in the infrastructure layer. The VMware connectors provide communication to the components that make up a vSAN cluster. SmartFabric Services provides automated network provisioning capabilities to the VxRail management platform. More details about SmartFabric Services are described in this [section](#). To provide serviceability features in VxRail, communication with ESRS and eServices need to be established. For all the different types of information VxRail gathers, a database is used to store metrics and a message queue is used to facilitate the necessary sequence of events/transactions for automation and orchestration of processes.

The local cluster management layer is the set of proprietary services built to provide customers the benefit of operational simplicity and intelligence to more effectively manage their VxRail cluster.

- System settings – hardware enclosure and components status and information are propagated to the VxRail
- SSO Authentication – integration with the vCenter Server single sign-on service
- Lifecycle management – end-to-end upgrade of hardware and software components of VxRail
- Marketplace – access to download VxRail ecosystem software from a single area
- Cluster operations – cluster expansion or node removal services
- Logging/audit – system logging and auditing services that has the capability to send data to vRealize Log Insight



- Health alerts – notifications of component health
- Serviceability – access to ESRS for customer support and eServices for product knowledge
- Data collector – collects cluster metrics and sends to VxRail ACE for infrastructure machine learning
- vCenter plugin – VxRail Manager interface for local cluster management
- First Run – automated deployment wizard for Professional Services to deliver VxRail on Day 0
- Network management – automated network provisioning services enabled by SmartFabric Services

Customers looking to leverage these services to manage their VxRail clusters locally can use the VxRail Manager plugin on their vCenter Server. There is also upstream APIs that provide a subset of these capabilities, such as lifecycle management and cluster operations services, to cloud solution developers looking to orchestrate the provisioning of HCI for their cloud service delivery solutions. These APIs are also beneficial to customers looking to script their own solutions to manage VxRail clusters at scale.

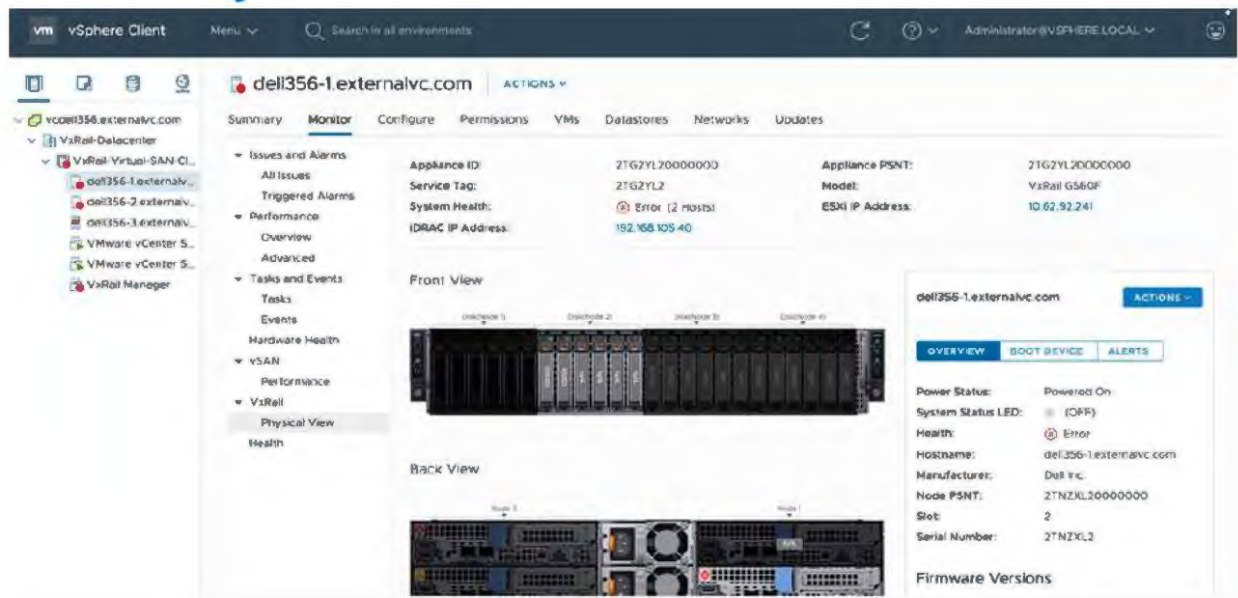
## VxRail Manager

VxRail Manager features user-friendly workflows for automating VxRail deployment and configuration and monitoring the health of individual systems and individual nodes in the entire cluster. It also incorporates functionality for hardware serviceability and system platform lifecycle management. For instance, it guides system administrators through adding new systems to an existing cluster, and it automatically detects new systems when they come online. VxRail Manager is also used to replace failed disk drives without disrupting availability, to generate and download diagnostic log bundles, and to apply VMware updates or software patches non-disruptively across VxRail nodes.

File-based backups of VxRail HCI System Software help to ensure business continuity in the rare event the VxRail VM needs to be rebuilt.

VxRail HCI System Software is preinstalled on the VxRail system as a single virtual machine. With VxRail Manager plugin for vCenter Server, all VxRail Manager features are integrated with and accessible from the vCenter Server so that users can benefit from these valuable capabilities on a familiar management interface. With the VxRail Manager plugin, the vCenter Server can manage physical hardware of the VxRail cluster.

Figure 17 VxRail Manager plugin for vCenter Server



VxRail also leverages VMware vRealize Log Insight to monitor system events and provide ongoing holistic notifications about the state of virtual environment and system hardware. It delivers real-time automated log management for the VxRail system with log monitoring, intelligent grouping, and analytics to provide better troubleshooting at scale across VxRail physical, virtual, and cloud environments. Furthermore, VxRail HCI System Software simplifies system platform lifecycle management by delivering patch software and update notifications that can be automatically installed without interruption or downtime.

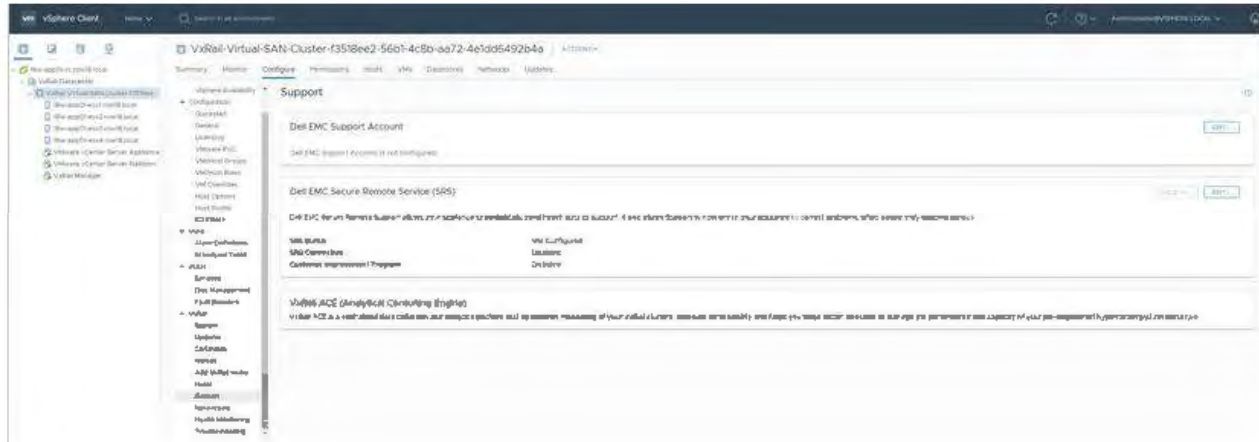
Dell EMC Software Remote Services (SRS), also accessible from within VxRail Manager plugin or REST API, provide enterprise-class support and services. SRS includes online chat support and Dell EMC field-service assistance (as seen in the figure below).

Figure 18 VxRail Manager ESRS details



In addition to SRS-specific support, the VxRail Support page on vCenter Server links to VxRail Community pages for Dell EMC Knowledge Base articles, user forums for FAQ information and VxRail best practices. The figure below is an example of the support view.

Figure 19 VxRail Manager Support Tab



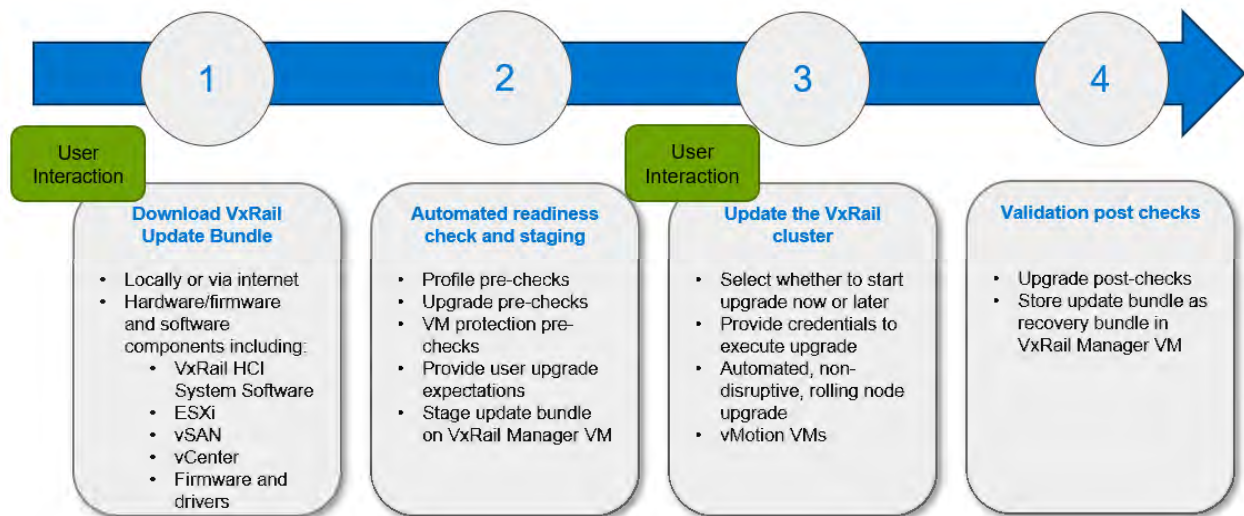
VxRail Manager plugin provides access to a digital market for finding and downloading qualified software packages such as VMware Horizon Cloud, Data Domain Virtual Edition, RecoverPoint for VM, vSphere Data Protection and other software options for VxRail systems.

### Customer upgradeable software

VxRail HCI system software is customer upgradeable via a fully automated and validated process. The single-click software upgrade is initiated from VxRail Manager plugin or REST API, and it automatically downloads all software ready to be updated including VxRail HCI System Software, vCenter Server and PSC, ESXi hosts, and SRS. The automated process consists of four steps including download of the VxRail software, a readiness check, the actual update of the software, and finally, validation and upgrade post checks. The final validation step ensures the upgrade was successful, and the VxRail system is fully functional at the new, upgraded version of software.

The figure below shows the four automated steps of a customer executed VxRail HCI system software upgrade. There are two touchpoints for the customer in this workflow, first when choosing where to acquire the update package and second to execute the cluster upgrade (hardware firmware and software update together)

Figure 20 Automated process steps for customer-executable VxRail HCI system software upgrade



Step 3 is performed one node at a time, where the ESXi host is placed in maintenance mode, and using vMotion, the VMs are moved to other nodes making the upgrade process non-disruptive. In the latest VxRail software versions, the upgrade process pre-stages the update bundle on the next node in the upgrade sequence while the current node is being upgraded. This improvement reduces the time to upgrade the node, ultimately reducing the overall time to complete a cluster upgrade.

### VxRail Analytical Consulting Engine (ACE)

VxRail ACE is a centralized data collection and analytics platform that streamlines the monitoring of multiple VxRail clusters for a customer, improves serviceability, and helps the customer make better decisions to manage performance and capacity of their HCI. It is a cloud-based analytics platform that leverages advanced telemetry collected from the VxRail clusters for its infrastructure machine learning to provide reporting and actionable insight. VxRail ACE's infrastructure machine learning utilizes built-in knowledge of Dell EMC best practices and more than 700 common issues. This product provides health scores for the entire HCI stack to enable customers to quickly identify areas to troubleshoot and to address areas to efficiently scale based on projected growth of IT resources.

VxRail ACE is available with no additional hardware or software required for the VxRail cluster. It relies on a data collector service provided by the VxRail HCI System Software running on the VxRail nodes to aggregate metrics from the vSAN cluster as well as from the VxRail system. Officially called the Adaptive Data Collector, the service transfers this data bundle to the VxRail ACE platform using the same Secure Remote Services (SRS) conduit for dial home services. Because it uses SRS, an Online Support account is required as well as configuration and enablement of SRS for data to be transferred to the VxRail ACE data lake. This repository is housed at Dell EMC. Using Pivotal Cloud Foundry as its cloud-based service platform, VxRail ACE incorporates its infrastructure machine learning to produce reporting and insight to enable customers to improve serviceability and operational efficiencies. VxRail ACE functionality is entirely consumed via a SaaS-based web portal which provides a single global view of the customer's VxRail environment.



## vSphere and vSAN ordering information

VxRail system allows customers to use any existing eligible vSphere licenses with their VxRail, or the licenses can be purchased with a VxRail. This VxRail vSphere license independent model (also called “bring your own” or BYO vSphere License model) allows customers to leverage a wide variety of vSphere licenses they may have already purchased.

Several vSphere license editions are supported with VxRail including Enterprise+, Standard, and ROBO editions (vSphere Enterprise is also supported, but is no longer available from VMware). Also supported are vSphere licenses from Horizon bundles or add-ons when the system is dedicated to VDI.

If vSphere licenses need to be purchased, they should be ordered through Dell EMC, the customer’s preferred VMware channel partner, or from VMware directly. Licenses acquired through VMware ELA, VMware partners, or Dell EMC will receive single-call support from Dell EMC.


When determining the best vSphere license to use with the VxRail system, a key consideration is the effect of VxRail functionality. DRS, a significant vSphere feature described earlier in this TechBook, provides the greatest amount of functional variance to VxRail clusters. Customers should consider the degree of automation that DRS provides to determine if the vSphere license they desire includes this functionality.

VxRail supports flexible vSAN licensing options and requires vSAN to be ordered with VxRail or applied via a vSAN ELA from VMware. VxRail supports all license editions of vSAN including Standard, Advanced, and Enterprise. For details on the differences between the vSAN versions please consult the VMware vSAN Comparison Guide here:

<https://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/products/vsan/vmware-vsan-67-licensing-guide.pdf>

The figure below lists the software included with VxRail.

Figure 21 VxRail included software



**Note:** Check with your Dell EMC representative to verify the latest software versions levels supported.

**Note:** The vCenter license included with VxRail is for vCenter deployed on VxRail and is not transferrable to another external server. All other software integrated with VxRail is also non-transferrable.

Use the information in the following figures as a guide.

Figure 22 General vSphere options



 <b>vSphere Enterprise Plus</b>	 <b>vSphere Standard</b>
<p><b>Dramatically increases administrator productivity</b></p> <ul style="list-style-type: none"> <li>• Automated workload rebalancing and affinity rules</li> <li>• Automated maintenance mode</li> <li>• Streamlined drive replacement</li> <li>• One-click software updates</li> <li>• vGPU support for VxRail V Series</li> </ul>	<p><b>Lower upfront costs; manual administration</b></p> <ul style="list-style-type: none"> <li>• Manual workload balancing</li> <li>• Manual maintenance mode</li> <li>• Multi-step drive replacement</li> <li>• One-click software updates</li> <li>• Does not support vGPU</li> </ul>

Figure 23 vSphere editions

VSPHERE EDITIONS & PACKAGES	VSPHERE EDITION FUNCTIONAL LEVEL	VSPHERE LICENSE KEY INPUT UNIT
vSphere Enterprise Plus	ENT+	CPU
vSphere Enterprise	ENT+	CPU
vSphere Standard	STD	CPU
vSphere ROBO Advanced	STD* (Additional Features, no DRS)	VM (Qty 25)
vSphere ROBO Standard	STD	VM (Qty 25)
vSphere Desktop	ENT+, virtual desktop only	User (Qty 100)
vSphere Platinum	ENT+	CPU + subscription
vSphere with Operations Management Enterprise Plus (VSOM)	ENT+	CPU
Horizon Enterprise	ENT+, virtual desktop only	User (Qty 10 or 100)
Horizon Advanced	ENT+, virtual desktop only	User (Qty 10 or 100)
Horizon Standard	ENT+, virtual desktop only	User (Qty 10 or 100)
vCloud Suite	ENT+	CPU

## VMware vSphere

The VMware vSphere software suite delivers an industry-leading virtualization platform to provide application virtualization within a highly available, resilient, efficient on-demand infrastructure—making it the ideal software foundation for VxRail systems. ESXi and vCenter are components of the vSphere software suite. ESXi is a hypervisor installed directly onto a physical VxRail server node, enabling it to be partitioned into multiple logical servers or virtual machines. Virtual machines are configured on top of the ESXi server. VMware vCenter server is a centralized management application that is used to manage the ESXi hosts and VMs.

The following sections provide an in-depth examination of the VMware software components as implemented in the VxRail software architecture.

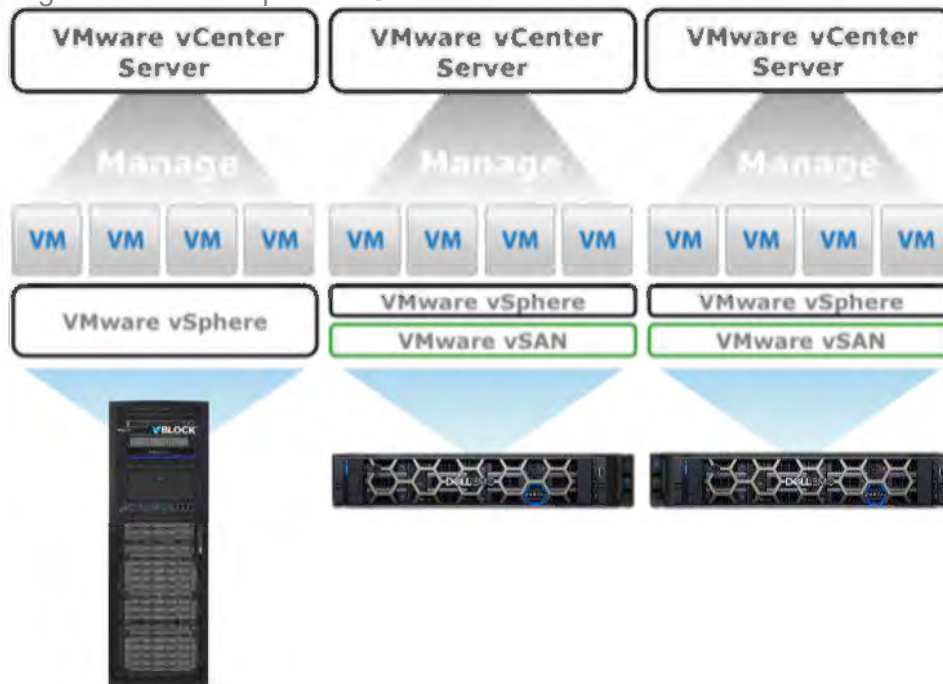
### VMware vCenter Server

vCenter Server is the centralized platform for managing a VMware environment. VxRail includes a license to run vCenter hosted on VxRail (the license is not transferrable to run vCenter on an external server). It is the primary point of management for both server virtualization and vSAN and is the enabling technology for advanced capabilities such as vMotion, Distributed Resource Scheduler (DRS), and high availability (HA). vCenter scales to enterprise levels where a single vCenter can support up to 1,000 hosts (VxRail nodes) and 10,000 virtual machines. vCenter supports a logical hierarchy of datacenters, clusters, and hosts, which allows resources to be segregated by use cases or lines of business and allows resources to move dynamically as needed. This is all done from a single interface.

As part of a VxRail deployment, a standalone vCenter instance is optionally configured on VxRail and provides the primary point of management for both the virtual machine and vSAN environments. The vCenter instance will be configured with an external Platform Services Controller (PSC), where the vCenter server and the PSC are configured as separate virtual machines.

The figure below is an example of three separate vCenter environments.

Figure 24 Three separate vCenter environments



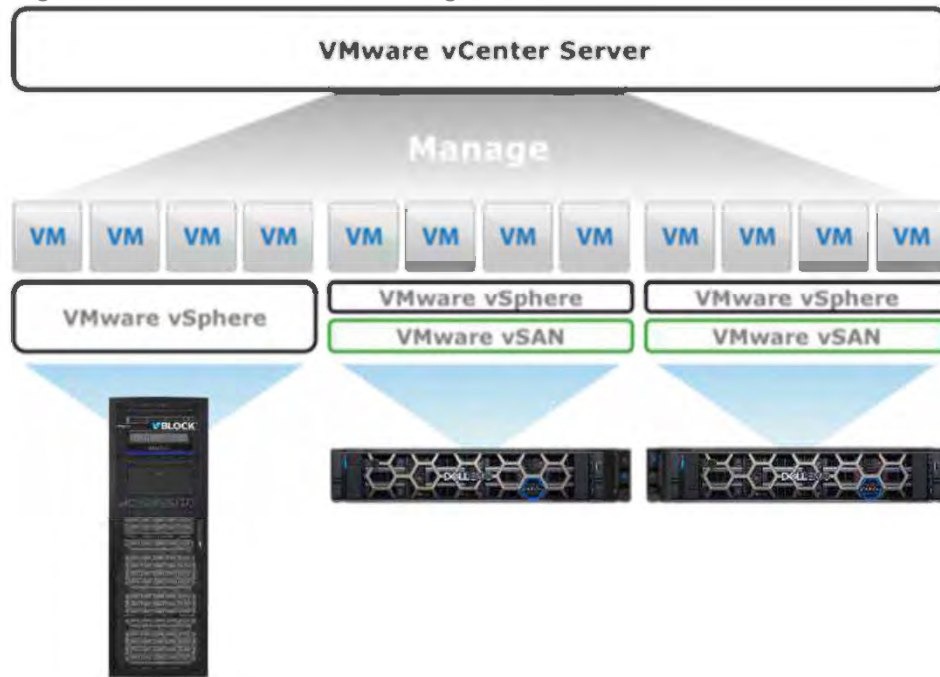
The VxBlock and two VxRail environments are managed as separate entities. This is a simple environment, and in some cases, the planning and deployment may be easier as there are no interactions.

A VxRail system can optionally join an existing externally hosted vCenter Server environment during the initial configuration (if the versions are compatible with the VxRail software). This allows for a central vCenter Server to manage multiple VxRail systems from a single pane of glass. Each VxRail cluster appears within vCenter.

The figure below shows an example where multiple VxRail clusters are part of an existing vCenter environment. Each VxRail is a unique and independent vSAN data store within vCenter.



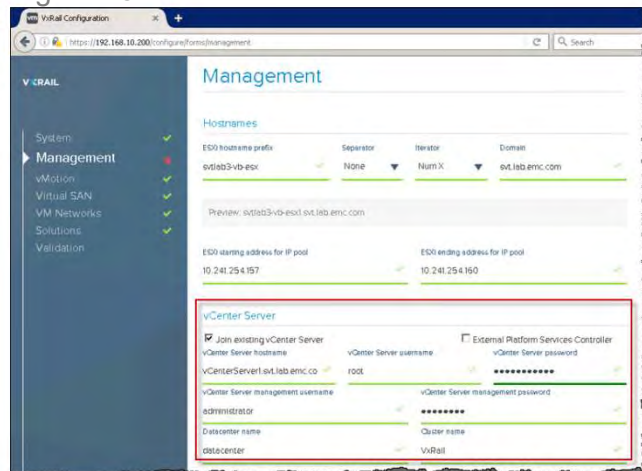
Figure 25 VxRail within an existing vCenter environment



The external vCenter server can be a physical server or a virtual server running as either a VCSA or in a Windows virtual machine. The Platform Services Controller (PSC) can be either embedded or non-embedded. As part of the initial configuration, the management page of the VxRail wizard presents the option to join an existing vCenter. If selected, specify the hostname of the vCenter Server and administrator password, identify the datacenter to add the VxRail environment, and supply the name of the cluster.

The figure below shows an example of the VxRail Initialization Wizard dialog to specify joining an existing vCenter Server.

Figure 26 Initialization wizard



The datacenter must already exist within vCenter, and the cluster will be created as part of the installation process. See the [VxRail vCenter Server Planning Guide](#) for more information on configuration options.

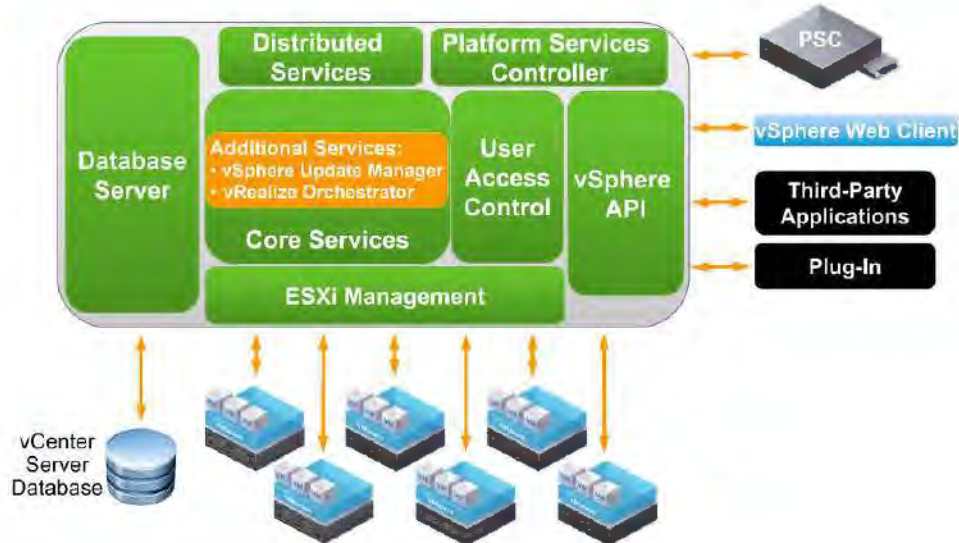
## vCenter services and interfaces

vCenter provides a number of services and interfaces, including:

- Core VM and resource services such as an inventory service, task scheduling, statistics logging, alarm and event management, and VM provisioning and configuration
- Distributed services such as vSphere vMotion, vSphere DRS, and vSphere HA
- vCenter Server database interface

The figure below clarifies the organization of vCenter services within the vSphere environment.

Figure 27 vCenter services



## PSC deployment options

The Platform Services Controller (PSC) can be deployed as either embedded (in externally hosted vCenter) or external, as shown below. When VxRail-hosts vCenter the PSC is deployed external as a separate VM. See the figure below.

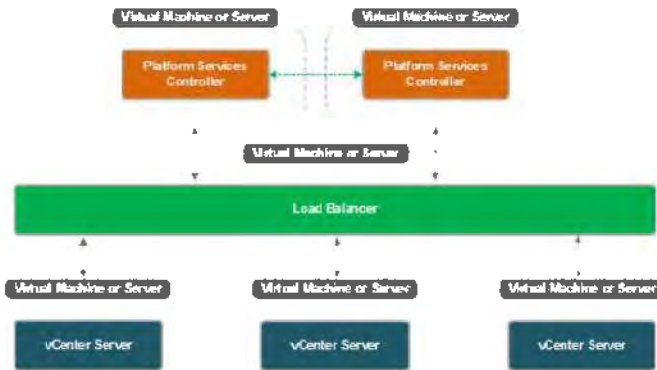
Figure 28 Embedded and external PCS deployments



Embedded PSC is ideal for small environments, or if simplicity and reduced resource utilization are key factors for the environment. The vCenter Server is bundled with an embedded PSC, and all the PSC services reside on the same host machine as vCenter Server.

External PSC (see the figure below) is ideal for larger environments, where there are multiple vCenter servers, but you want a single pane-of-glass for the site. The services bundled with the PSC and vCenter Server are deployed on different virtual machines or even different physical servers. When utilizing an externally hosted vCenter with external PSC, Enhanced Linked Mode (ELM) is supported.

Figure 29 External PSCs configured for high availability



## Enhanced Linked Mode

Multiple vCenter Servers can connect to the same external Platform Service Controller. Enhanced Linked Mode enables a consolidated management view of multiple vCenter Servers configured to use the Platform Services Controller domain. This includes a common inventory where an administrator can search for objects across vCenter Servers. Roles, permissions, licenses, and other key data across systems are also replicated across vCenter instances.

This single pane-of-glass view provides enterprise-level scale and works particularly well in large, multiple VxRail cluster environments or when the VxRail system joins an existing large vSphere environment. In addition, Enhanced Link Mode also enables capabilities such as Cross vCenter vMotion where VMs can be moved between vCenters.


Enhanced Linked Mode is supported with VxRail when running an externally hosted vCenter with VxRail and utilizing external PSC.

## VMware vSphere ESXi

vSphere is the core operational software in the VxRail system. vSphere aggregates a comprehensive set of features that efficiently pools and manages the resources available under the ESXi hosts. Keep in mind that this TechBook focuses on vSphere technology specifically as it pertains to the VxRail system. Features included in other vSphere implementations may not apply to VxRail and features included in VxRail may not apply to other implementations.

VMware ESXi is an enterprise-class hypervisor that deploys and services virtual machines. The following figure illustrates its basic architecture.

Figure 30 Birds-eye view: vSphere ESXi architecture



ESXi partitions a physical server into multiple secure and portable VMs that can run side by side on the same physical server. Each VM represents a complete system—with processors, memory, networking, storage, and BIOS—so any operating system (guest OS) and software applications can be installed and run in the virtual machine without any modification.


The hypervisor provides physical-hardware resources dynamically to virtual machines (VMs) as needed to support the operation of the VMs. The hypervisor enables virtual machines to operate with a degree of independence from the underlying physical hardware. For example, a virtual machine can be moved from one physical host to another. Also, the VM's virtual disks can be moved from one type of storage to another without affecting the functioning of the virtual machine.

ESXi also isolates VMs from one another, so when a guest operating system running in one VM fails, other VMs on the same physical host are unaffected and continue to run. Virtual machines share access to CPUs, and the hypervisor is responsible for CPU scheduling. In addition, ESXi assigns VMs a region of usable memory and provides shared access to the physical network cards and disk controllers associated with the physical host. Different virtual machines can run different operating systems and applications on the same physical computer.

## Communication between vCenter Server and ESXi hosts

vCenter Server communicates with the ESXi host through a vCenter Server agent, also referred to as vpxa or the vmware-vpxa service, which is started on the ESXi host when it is added to the vCenter Server inventory. See the figure below.

Figure 31 Communication between vCenter and ESXi hosts



Specifically, the vCenter vpxd daemon communicates through the vpxa service to the ESXi host daemon known as the hostd process. The vpxa process acts as an intermediary between the vpxd process that runs on vCenter Server and the hostd process that runs on the ESXi host, relaying the tasks to perform on the host. The hostd process runs directly on the ESXi host and is responsible for managing most of the operations on the ESXi host including creating VMs, migrating VMs, and powering on VMs.

## Virtual machines

A virtual machine consists of a core set of the following related files, or a set of objects, as shown in the figure below.

Figure 32 Virtual machine files

Name	Type	Size
.dvsData	Folder	
sdd sf	Folder	
Test.VM1-572c3eef.v...	File	1,060,449.23 KB
Test.VM1.nvram	Non-volatile Memory File	848 KB
Test.VM1.vmdk	Virtual Disk	6,711,296.03 KB
Test.VM1.vmsd	File	0.03 KB
Test.VM1.vmx	Virtual Machine	3.44 KB
vmware-1.log	VM Log File	374.14 KB
vmware-2.log	VM Log File	221.93 KB
vmware-3.log	VM Log File	219.72 KB
vmware.lcg	VM Log File	218.44 KB

Except for the log files, the name of each file starts with the virtual machine's name (VM\_name). These files include:

A configuration file (.vmx) and/or a virtual-machine template-configuration file (.vmtx)

One or more virtual disk files (.vmdk)

A file containing the virtual machine's BIOS settings (.nvram)

A virtual machine's current log file (.log) and a set of files used to archive old log entries (-#.log)

Swap files (.vswp), used to reclaim memory during periods of contention




A snapshot description file (.vmsd), which is empty if the virtual machine has no snapshots

## Virtual machine hardware

A virtual machine uses virtual hardware. Each guest operating system sees ordinary hardware devices and does not know that these devices are virtual. Hardware resources are shown in the figure below.

Figure 33 Hardware resources for VMs



All virtual machines have uniform hardware, except for a few variations that the system administrator can apply. Uniform hardware makes virtual machines portable across VMware virtualization platforms. vSphere supports many of the latest CPU features, including virtual CPU performance counters. It is possible to add virtual hard disks and NICs, and configure virtual hardware, such as CD/DVD drives, floppy drives, SCSI devices, USB devices, and up to 16 PCI vSphere DirectPath I/O devices.

## Virtual Machine Communication

The Virtual Machine Communication Interface (VMCI) provides a high-speed communication channel between a virtual machine and the hypervisor. VMCI devices cannot be added or removed. The SATA controller provides access to virtual disks and DVD/CD-ROM devices. The SATA virtual controller appears to a virtual machine as an AHCI SATA controller. Without VMCI, virtual machines would communicate with the host using the network layer, which adds overhead to the communication. With VMCI, communication overhead is minimal, and tasks requiring that communication can be optimized. An internal network can transmit an average of slightly over 2Gbps using VMXNET3. VMCI can go up to nearly 10Gbps with twelve 8k-sized queue pairs.

VMCI provides socket APIs that are very similar to the APIs already used for TCP/UDP applications.

For more information about the virtual hardware, see *vSphere Virtual Machine Administration Guide* at <https://www.vmware.com/support/pubs/vsphere-esxi-vcenter-server-6-pubs.html>.

## Virtual networking

VMware vSphere provides a rich set of networking capabilities that integrate well with sophisticated enterprise networks. These networking capabilities are provided by the ESXi server and managed by vCenter. Virtual networking provides the ability to network virtual machines in the same way physical machines are networked. Virtual networks can be built within a single ESXi server host or across multiple ESXi server hosts. VxRail ESXi hosts use a virtual switch for communication among virtual machines in the VxRail cluster using the same protocols that would be used over physical switches, without the need for additional networking hardware. The virtual switch also supports VLANs that are compatible with standard VLAN implementations from switch vendors. A virtual switch, like a physical Ethernet switch, forwards frames at the data-link layer.

Virtual Ethernet adapters are the key vSphere components for virtual networking. A virtual machine can be configured with one or more virtual Ethernet adapters, each of which has its own IP address and MAC address. As a result, virtual machines have the same properties as physical machines from a networking standpoint. In addition, virtual networks enable functionality not possible with physical networks today. Virtual Ethernet adapters are used by individual virtual machines and the virtual switches that connect VMs to each other and connect both virtual machines and the ESX Server service console to external networks.

The virtual switch links to the external network through outbound Ethernet adapters called vmnics, and the virtual switch can bind multiple vmnics together (much like NIC teaming on a traditional server), extending availability and bandwidth to the virtual machines it services.


Virtual switches are similar to their physical-switch counterparts. Like a physical network device, each virtual switch is isolated for security and has its own forwarding table. An entry in one table cannot point to another port on another virtual switch. The switch looks up only destinations that match the ports on the virtual switch where the frame originated. This feature stops potential hackers from breaking virtual switch isolation. Virtual switches also support VLAN segmentation at the port level, so each port can be configured either as an access port to a single VLAN or as a trunk port to multiple VLANs.

### Virtual Distributed Switch

VxRail clusters use the VMware Virtual Distributed Switch (VDS), which functions as a single switch that spans across multiple nodes in the same cluster. This switch enables virtual machines to maintain consistent network configuration as they migrate across multiple hosts. A distributed switch is configured in vCenter Server at the datacenter level and makes the configuration consistent across all hosts. vCenter Server stores the state of distributed ports in the vCenter Server database. Networking statistics and policies migrate with virtual machines when the virtual machines are moved from host to host. As discussed in upcoming sections, vSAN relies on VDS for its storage-virtualization capabilities, and the VxRail system uses VDS for system traffic.

The following figure provides an overview of VDS.

Figure 34 Virtual Distributed Switch



## vMotion and Virtual Machine mobility

VMware vMotion enables the live migration of running virtual machines from one physical server to another with zero downtime, continuous service availability, and complete transaction integrity. vMotion is a key enabling technology for creating the dynamic, automated, and self-optimizing datacenter. vMotion continuously and automatically allocates virtual machines within resource pools. It also improves availability by conducting maintenance without disrupting business operations.

The advanced capability for migrating workloads without disruption is one of the features that distinguish the VxRail solution from other HCI options. In the vSphere virtual infrastructure, migration refers to moving a virtual machine from one host, datastore, or vCenter Server system to another host, datastore, or vCenter Server system. Different types of migrations exist including:

Cold, which is migrating a powered-off VM to a new host or datastore

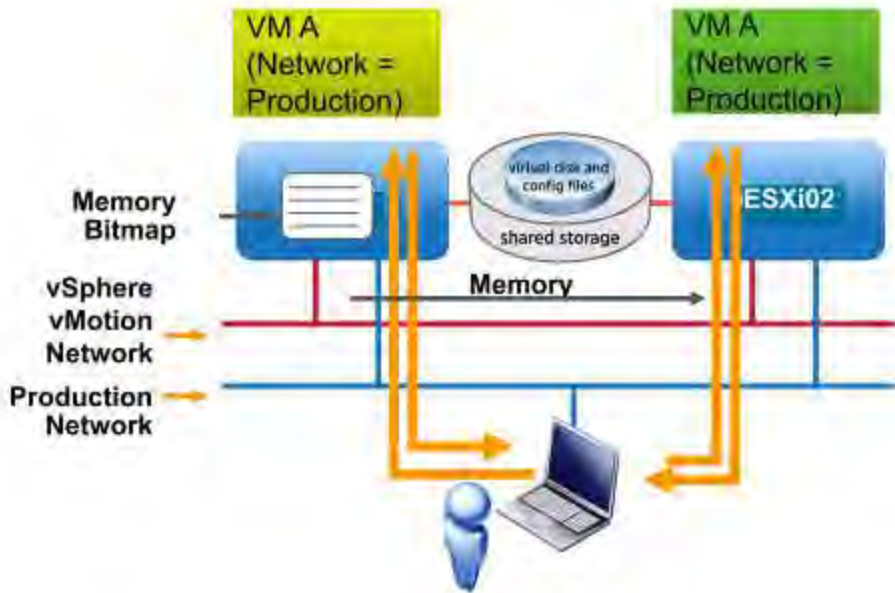
Suspended, which is migrating a suspended VM to a new host or datastore

Live, which uses vSphere vMotion to migrate a “live,” powered-on VM to a new host and/or uses vSphere Storage vMotion to migrate the files of a live, powered-on VM to a new datastore



vMotion allows for live migration of virtual machines between ESXi hosts without disruption or downtime. The process is summarized in the figure below.

Figure 35 vMotion migration



With vMotion, while the entire state of the virtual machine is migrated, the data remains in the same datastore. The state information includes the current memory content and all the information that defines and identifies the virtual machine. The memory content consists of transaction data and whatever bits of the operating system and applications in memory. The definition and identification information stored in the state includes all the data that maps to the virtual machine hardware elements, including BIOS, devices, CPU, and MAC addresses for the Ethernet cards.

A vMotion migration consists of the following steps:

1. The VM memory state is copied over the vMotion network from the source host to the target host. Users continue to access the VM and, potentially, update pages in memory. A list of modified pages in memory is kept in a memory bitmap on the source host.
2. After most of the VM memory is copied from the source host to the target host, the VM is quiesced. No additional activity occurs on the VM. During the quiesce period, vMotion transfers the VM-device state and memory bitmap to the destination host.
3. Immediately after the VM is quiesced on the source host, the VM is initialized and starts running on the target host. A Gratuitous Address Resolution Protocol (GARP) request notifies the subnet that the MAC address for the VM is now on a new switch port.
4. Users access the VM on the target host instead of the source host. The memory pages used by the VM on the source host are marked as free.

## Enhanced vMotion Compatibility

Enhanced vMotion Compatibility (EVC) is a cluster feature that prevents vMotion migrations from failing because of incompatible CPUs. EVC ensures that all hosts in a cluster present the same CPU feature set to virtual machines, even if the actual CPUs on the hosts differ. It prevents migration failures due to CPU incompatibility. This is on by default in VxRail systems.


## Storage vMotion

Storage vMotion uses an I/O-mirroring architecture to copy disk blocks between source and destination:

1. Initiate storage migration.
2. Use the VMkernel data mover and provide vSphere Storage APIs for Array Integration (VAAI) to copy data.
3. Start a new VM process.
4. Mirror I/O calls to file blocks that have already been copied to virtual disk on the target datastore. Switch to the target-VM process to begin accessing the virtual-disk copy.

The figure below illustrates the process

Figure 36 Storage vMotion



The storage-migration process copies the disk just once, and the mirror driver synchronizes the source and target blocks with no need for recursive passes. In other words, if the source block changes after it migrates, the mirror driver writes to both disks simultaneously which maintains transactional integrity. The mirroring architecture of Storage vMotion produces more predictable results, shorter migration times, and fewer I/O operations than more conventional storage-migration options. It's fast enough to be unnoticeable to the end user. It also guarantees migration success even when using a slow destination disk.

vSphere supports the following Storage vMotion migrations:

- Between clusters
- Between datastores (including non-vSAN to vSAN and vice versa)
- Between networks
- Between vCenter Server instances for vCenter Servers configured in Enhanced Link Mode with hosts that are time-synchronized
- Over long distances (up to 150ms round-trip time)

Note that for VxRail clusters, Storage vMotion can only be used for migration into or out of the vSAN datastore.

## vSphere Distributed Resource Scheduler

VMware Distributed Resource Scheduler (DRS) is a key feature included with vSphere Enterprise Plus and vSphere with Operations Management Enterprise Plus. DRS balances computing capacity across a collection of VxRail server resources that have been aggregated into logical pools. It continuously balances and optimizes compute resource allocation among the VMs.

When a VM experiences an increased workload, DRS evaluates the VM priority against user-defined resource-allocation rules and policies. If justified, DRS allocates additional resources. It can also be configured to dedicate consistent resources to the VMs of particular business-unit applications to meet SLAs and business requirements.

DRS allocates resources to the VM either by migrating the VM to another server with more available resources or by making more “resources” for the VM on the same server by migrating other VMs off the server. In the VxRail system, all ESXi hosts are part of a vMotion network. The live migration of VMs to different node servers is completely transparent to end users through vMotion (see the figures below). DRS adds tremendous value to the VxRail cluster by automating VM placement and ensuring consistent and predictable application-workload performance.

Figure 37 DRS movement of VMs across node servers


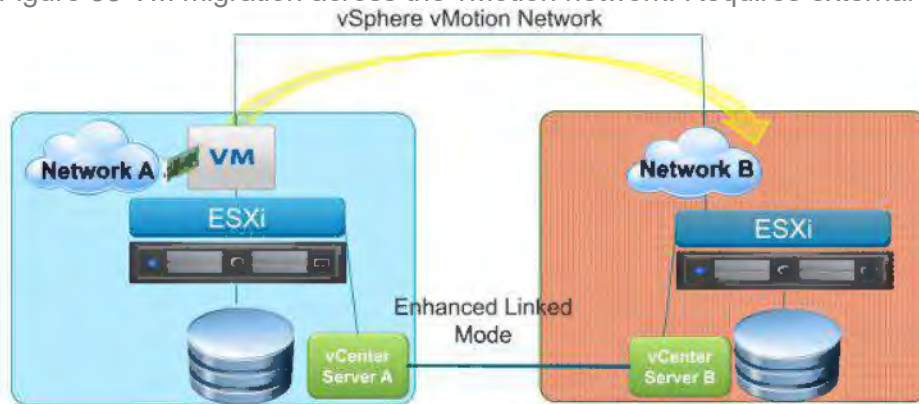


Figure 38 VM migration across the vMotion network. Requires external vCenter.

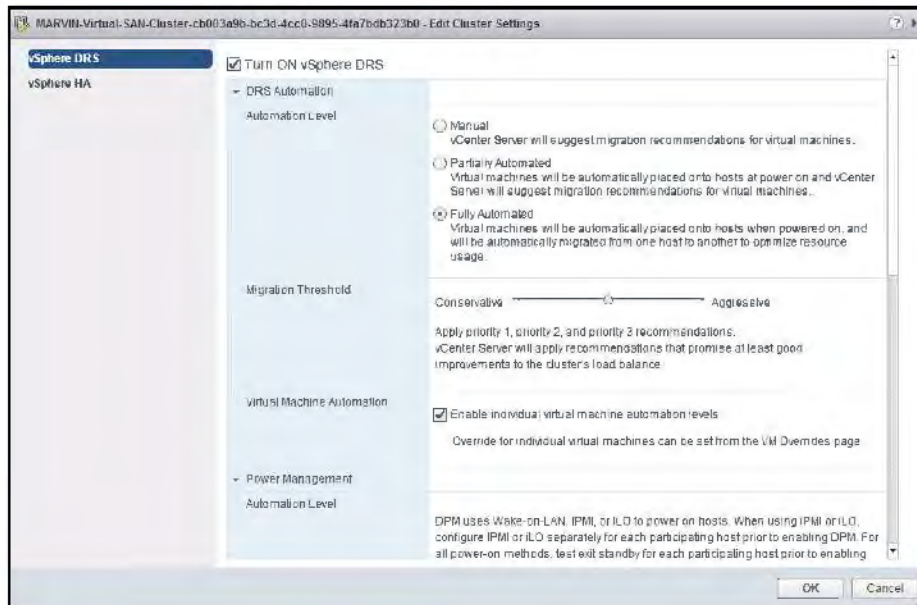


DRS offers a considerable advantage to VxRail users during maintenance situations, because it automates the tasks normally involved in manually moving live machines during upgrades or repairs. DRS facilitates maintenance automation, providing transparent, continuous operations by dynamically migrating all VMs to other physical servers. That way, servers can be attended to for maintenance, or new node servers can be added to a resource pool, all while DRS automatically redistributes the VMs among the available servers as the physical resources change.

In other words, DRS dynamically balances VMs as soon as additional resources become available when new server is added or when an existing server has finished its maintenance cycle. DRS allocates only CPU and memory resources for the VMs and uses vSAN for shared storage.

The following figure shows the settings for configuring DRS.

Figure 39 Configuring DRS settings



Some conditions and business operations warrant a more aggressive DRS migration strategy than others. Adjustable cluster parameters establish the thresholds that trigger DRS migrations (as shown in the screen shot above.) For example, a Level-2 threshold only applies specified migration recommendations to make a significant impact on the cluster's load balance, whereas a Level-5 threshold applies all the recommendations to even slightly improve the cluster's load balance.

DRS applies only to VxRail virtual machines. (vSAN uses a single datastore and handles placement and balancing internally. vSAN does not currently support Storage DRS or Storage I/O Control.)


## vSphere High Availability (HA)

vSphere provides several solutions to ensure a high level of availability, during both planned and unplanned downtime scenarios. vSphere depends on the following technologies to make sure that virtual machines running in the environment remain available:

- Virtual machine migration
- Multiple I/O adapter paths
- Virtual machine load balancing
- Failure tolerance
- Disaster recovery

Together with vSAN, vSphere HA produces a resilient, highly available solution for VxRail virtual machine workloads. vSphere HA protects virtual machines by restarting them in the event of a host failure. (See the figure below.) It leverages the ESXi cluster configuration to ensure rapid recovery from outages, providing cost-effective high availability for applications running in virtual machines. When a host joins a cluster, its resources become part of the cluster resources. The cluster manages the resources of all hosts within it. In a vSphere environment, ESXi clusters are responsible for vSphere HA, DRS, and the vSAN technology that provides VxRail software-defined storage capabilities. See the figure below.

Figure 40 vSphere HA





vSphere HA provides several points of protection for applications:

It circumvents any server failure by restarting the virtual machines on other hosts within the cluster.

It continuously monitors virtual machines and resets any detected VM failures.

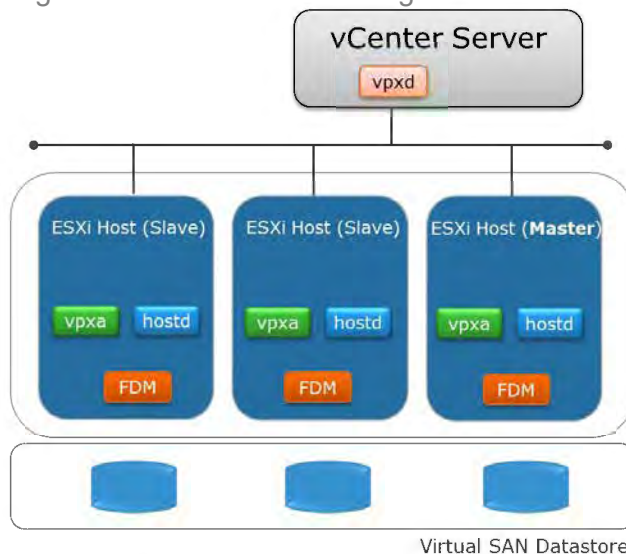
It protects against datastore accessibility failures and provides automated recovery for affected virtual machines. With Virtual Machine Component Protection (VMCP), the affected VMs are restarted on other hosts that still have access to the datastores.

It protects virtual machines against network isolation by restarting them if their host becomes isolated on the management or VMware vSAN network. This protection is provided even if the network has become partitioned.

Once vSphere HA is configured, all workloads are protected. No actions are required to protect new virtual machines and no special software needs to exist within the application or virtual machine.

Included in the failover capabilities in vSphere HA is a service called the Fault Domain Manager (FDM) that runs on the member hosts (shown in the figure below). After the FDM agents have started, the cluster hosts become part a fault domain, and a host can exist in only one fault domain at a time. Hosts cannot participate in a fault domain if they are in maintenance mode, standby mode, or disconnected from vCenter Server.

Figure 41 Fault Domain Manager



vpxa: Intermediary between hostd and vpxd  
vpxd: vCenter agent that communicates with hostd  
hostd: responsible for managing ESXi operations  
FDM: Fault Domain Manager

FDM uses a master-slave operational model (see the figure above). An automatically designated master host manages the fault domain, and the remaining hosts are slaves. FDM agents on slave hosts communicate with the FDM service on the master host using a secure TCP connection. In VxRail clusters, vSphere HA is enabled only after the vSAN cluster has been configured. Once vSphere HA has started, vCenter Server contacts the master host agent and sends it a list of cluster-member hosts along with the cluster configuration. That information is saved to local storage on the master host and then pushed out to the slave hosts in the cluster. If additional hosts are added to the cluster during normal operation, the master agent sends an update to all hosts in the cluster.

The master host provides an interface to vCenter Server for querying and reporting on the state of the fault domain and virtual-machine availability. vCenter Server governs the vSphere HA agent, identifying the virtual machines to protect and maintaining a VM-to-host compatibility list.

The agent learns of state changes through hostd, and vCenter Server learns of them through vpxa. The master host monitors the health of the slaves and takes responsibility for virtual machines that had been running on a failed slave host. Meanwhile, the slave host monitors the health of its local virtual machines and sends state changes to the master host. A slave host also monitors the health of the master host.

vSphere HA is configured, managed, and monitored through vCenter Server. Cluster configuration data is maintained by the vCenter Server vpxd process. If vpxd reports any cluster configuration changes to the master agent, the master advertises a new copy of the cluster configuration information and then each slave fetches the updated copy and writes the new information to local storage. Each datastore includes a list of protected virtual machines. The list is updated after vCenter Server notices any user-initiated power-on (protected) or power-off (unprotected) operation.


## vCenter Server Watchdog

One method of providing vCenter Server availability is to use the Watchdog feature in a vSphere HA cluster. Watchdog monitors and protects vCenter Server services. If any services fail, Watchdog attempts to restart them. If it cannot restart the service because of a host failure, vSphere HA restarts the virtual machine running the service on a new host. Watchdog can provide better availability by using vCenter Server processes (PID Watchdog) or the vCenter Server API (API Watchdog).

## vSphere Encryption

vSphere encryption enables customers to encrypt data on a per VM level. This level of encryption is ideal for customers who are concerned about rogue admins sending a VM and all its data to a non-secure location. Which VMs should be encrypted, is up to the virtualization administrative team and can be selected on a per VM basis (as seen in the figure below). A KMIP-compliant Key Management Server like CloudLink or Hytrust is required.

Figure 42 Per VM-level encryption with vSphere Encryption



## vSAN


VxRail systems leverage VMware's vSAN software, which is fully integrated with vSphere and provides full-featured and cost-effective software-defined storage. vSAN implements a notably efficient architecture, built directly into hypervisor. This distinguishes vSAN from solutions that typically install a virtual storage appliance (VSA) that runs as a guest VM on each host. Embedding vSAN into the ESXi kernel layer has obvious advantages in performance and memory requirements. It has very little impact on CPU utilization (less than 10 percent) and self-balances based on workload and resource availability. It presents storage a familiar datastore construct and works seamlessly with other vSphere features such as vMotion.

vSAN aggregates locally attached disks of hosts in a vSphere cluster to create a pool of distributed shared storage. Capacity is easily scaled up by adding additional disks to hosts in the cluster and scaled out by adding additional ESXi hosts. This provides the flexibility to start with a very small environment and scale it over time. Storage characteristics are configured using Storage Policy Based Management (SPBM), which allows object-level policies to be set and modified on the fly to control storage provisioning and day-to-day management of storage service-level agreements (SLAs). vSphere and vSAN provide the foundation for VxRail performance and scale, and enable enterprise-class capabilities for hosted virtual machines including mobility using vMotion, High Availability, and Dynamic Resource Scheduler (DRS).

vSAN is preconfigured when the VxRail system is first initialized and managed through vCenter. The VxRail system-initialization process discovers locally attached storage disks from each ESXi node in the cluster to create a distributed, shared-storage datastore. The amount of storage in the vSAN datastore is an aggregate of all of the capacity drives in the cluster.

The figure below shows an example of a hybrid configuration where each node contributes storage capacity to the vSAN datastore. The SSD drive provides caching to optimize performance and hard disk drives (HDD) for capacity.

Figure 43 vSAN datastore



vSAN enables rapid storage provisioning within vCenter as part of the VM-creation and deployment operations. vSAN is policy driven and designed to simplify storage provisioning and



management. It automatically and dynamically matches requirements with underlying storage resources based on VM-level storage policies. VxRail provides two different vSAN node-storage configuration options: a hybrid configuration that leverages both flash SSDs and mechanical HDDs, and an all-flash SSD configuration. The hybrid configuration uses flash SSDs at the cache tier and mechanical HDDs for capacity and persistent data storage. The all-flash configuration uses flash SSDs for both the caching tier and capacity tier.

## Disk groups

Disk drives in VxRail hosts are organized into disk groups, and disk groups contribute storage to the vSAN cluster. In a VxRail system, a disk group contains a maximum of one flash-cache drive and six capacity devices. Depending on the model, each VxRail node can be configured with up to four disk groups. The following figure shows the number of disk groups per node and the number of drives per disk group for each VxRail model.

Figure 44 Disk group configurations for VxRail based on 14th generation servers

Model	Disk Slots	Fixed Cache Disk Slots	Maximum Disk Groups (DG)	Max Capacity Disks Per DG
E Series	10 (2.5-inch)	8,9	2	4
P Series, V Series	24 (2.5-inch)	20,21,22,23	4	5
S Series	12 (3.5-inch) + 2 (2.5-inch)	12,13 (2.5-inch)	2	6

The disk group layout, based on capacity, performance, and availability, is determined when the system is designed and sized. Below are the considerations for designing disk-group layout:

**The number of SSD cache drives.** Each disk group requires one and only one high-endurance SSD flash-cache drive. More disk groups require more SSD flash-cache drives. Currently, the VxRail system is offered with 400GB, 800GB, and 1600GB cache flash devices. For VxRail, the SSD cache drive must be installed in designated slots. See product documentation for specific slot locations.

**The total number of capacity drives available.** Each disk group requires at least one capacity drive. Only capacity drive counts when determining the total capacity available for a vSAN Cluster.

**Cache requirements.** Cache improves read and write performance in a hybrid storage configuration and write performance and endurance in an all-flash configuration. The optimal amount of cache depends on the active working set size. Workloads that have a larger working set size may require more cache, and this can be configured by either using larger SSD cache drives or configuring multiple disk groups with each disk group having one cache drive.

**Note:** All-flash configurations support a maximum of 600GB of write cache per disk group. Consider multiple disk groups for larger write cache per node.

**Performance.** For the same total capacity, a larger number of smaller drives will provide more IOPs than fewer larger drives. Depending on the model, configuring more drives may require multiple disk groups.

**Fault Domains and recoverability.** If the SSD cache drive fails, the full disk group goes offline and must be recovered. Larger disk groups take longer to recover. A node with only a single disk group will not be contributing any IOPs to the cluster during recovery, and this can reduce cluster performance.

## Hybrid and All-Flash differences

Cache is used differently in hybrid and all-flash configurations. In hybrid disk-group configurations (which use mechanical HDDs for capacity and flash SSD devices for the caching), the caching algorithm attempts to maximize both read and write performance. The flash SSD device serves two purposes: a read cache and a write buffer. 70 percent of the available cache is allocated for storing frequently read disk blocks, minimizing accesses to the slower mechanical disks. The remaining 30 percent of available cache is allocated to writes. Multiple writes are coalesced and written sequentially if possible, again maximizing mechanical HDD performance.

In all-flash configurations, one designated SSD flash device is used for the cache, while additional SSD flash devices are used for the capacity. In all-flash disk-group configurations, there are two types of flash SSDs: a very fast and durable flash device that functions as write cache and more cost-effective SSD devices that function as capacity. Here, the cache-tier SSD is 100 percent allocated for writes. None of the flash cache is used for reads; read performance from capacity-tier flash SSDs is more than sufficient for high performance. Many more writes can be held by the cache SSD in an all-flash configuration, and writes are only written to capacity when needed, which extends the life of the capacity-tier SSD.

While both configurations dramatically improve the performance of VMs running on vSAN, all-flash configurations provide the most predictable and uniform performance regardless of workload.

### Read cache: Basic function

The read cache, which only exists in hybrid configurations, keeps a collection of recently read disk blocks. This reduces the I/O read latency in the event of a cache hit; that is, the disk block can be fetched from cache rather than mechanical disk. For a given VM data block, vSAN always reads from the same replica/mirror. When there are multiple replicas (to tolerate failures), vSAN divides up the caching of the data blocks evenly between the replica copies.

If the data block being read from the first replica is not in cache, the directory service is referenced to discover whether or not the data block exists in the cache of another mirror (on another host) in the cluster. If the data block is found there, the data is retrieved. If the data block is not in cache on the other host, then there is a read-cache miss. In that case, the data is retrieved directly from the mechanical HDD.

### Write cache: Basic function

The write cache, found in both hybrid and all-flash configurations, behaves as a non-volatile write buffer. This greatly improves performance in both hybrid and all-flash configurations and also extends the life of flash capacity devices in all-flash configurations. When writes are written to cache, vSAN ensures that a copy of the data is written elsewhere in the cluster. All VMs deployed with vSAN are set with a default availability policy that ensures at least one additional copy of the VM data is available. This includes making sure that writes end up in multiple write caches in the cluster.

Once an application running inside the guest OS initiates a write, it is duplicated to the write cache on the hosts that include replicas of the storage objects. This means that, in the event of a host failure, a copy of the data is in cache and no data loss occurs. The VM simply uses the replicated copy of the cache data.

## Flash endurance

Flash endurance is related to the number of write/erase cycles that the cache-tier flash SSD can tolerate over its lifespan. It is important to choose the right level of endurance for each application or use case, according to the level of write activity that is expected.

For vSAN and VxRail configurations, the endurance specification uses Terabytes Written per day (TBW) as a metric to determine endurance.

For all-flash vSAN deployments, the specification for the cache device is 4TBW per day, which is an appropriate endurance for write intensive workloads.

## vSAN impact on flash endurance

There are two commonly used approaches to improve NAND flash endurance: improve wear leveling and minimize write activity. A distributed storage implementation that focuses on localizing data on the same node where the VMs reside prevents the distribution of the writes across all the drives in the cluster. This localization inevitably increases drive usage, leading to early drive replacement.

In contrast, vSAN distributes the objects and components of a VM across all the disk groups in the VxRail cluster. This distribution significantly improves wear leveling and reduces write activity by deferring writes. vSAN also reduces writes by employing data-reduction techniques such as de-duplication and compression.

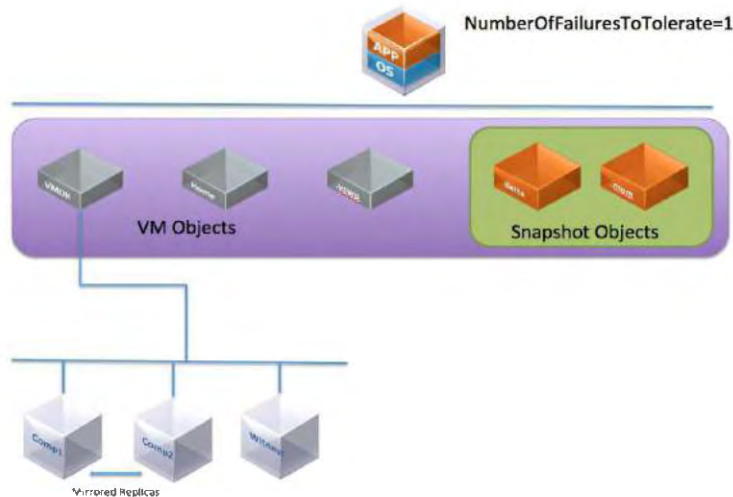
## Client cache

The client cache is used on both hybrid and all-flash configurations. It leverages local DRAM server memory (client cache) within the node to accelerate read performance. The amount of memory allocated is .4% -1GB per host. vSAN first tries to fulfill the read request from the local client cache, so the VM can avoid crossing the network to complete the read. If the data is unavailable in the client cache, the cache-tier SSD is queried to fulfill the read request. The client cache benefits read cache-friendly workloads.

## Objects and components

VxRail virtual machines are made up of a set of objects. For example, a VMDK is an object, a snapshot is an object, VM swap space is an object, and the VM home namespace (where the .vmx file, log files, and so on are stored) is also an object. See the figure below.

Figure 45 vSAN objects and components



Virtual-machine objects are split into multiple components based on performance and availability requirements defined in the storage policy applied to the objects of the VM. For example, if the VM is deployed with a policy to tolerate a single failure, the objects have two replica components. Distributed storage uses a disk-striping process to distribute data blocks across multiple devices. The stripe itself refers to a slice of data; the striped device is the individual drive that holds the stripe. If the policy contains a stripe width, the object is striped across multiple devices in the capacity layer, and each stripe is an object component.

Each vSAN host has a maximum of 9,000 components. The largest component size is 255GB. For objects greater than 255GB, vSAN automatically divides them into multiple components. For example, a VMDK of 62TB generates more than 500 x 255GB components. The figure above illustrates how components that make up VM objects are spread across drive on nodes based on Failure to tolerate policy.

## Witness

In vSAN, witnesses are an integral component of every storage object when the object is configured to tolerate at least one failure and when using mirroring as the Failure Tolerance Method (FTM). Witnesses are components that contain no data, only metadata. Their purpose is to serve as tiebreakers when availability decisions are made to meet the Failures to tolerate (FTT) policy setting, and they're used when determining if a quorum of components exist in the cluster.

In vSAN, storage components can be distributed in such a way that they can guarantee availability without relying on a witness. In this case, each component has a number of votes—at least one or more. Quorum is calculated based on the rule that requires "more than 50 percent of votes."

## Replicas


Replicas make up the virtual machine's storage objects. Replicas are instantiated when an availability policy (FTT) is specified for the virtual machine. The availability policy dictates how many replicas are created and lets virtual machines continue running with a full complement of data even when host, network, or disk failures occur in the cluster.

## Storage Policy Based Management (SPBM)

vSAN policies define virtual-machine storage requirements for performance and availability. These policies determine how storage objects are provisioned and allocated within the datastore to guarantee the required level of service.

vSAN implements Storage Policy Based Management, and each virtual machine deployed in a vSAN datastore has at least one assigned policy. When the VM is created and assigned a storage policy, the policy requirements are pushed to the vSAN layer. See the figure below.

Figure 46 SPBM



Storage policies are a set of rules and are assigned to VMs either manually or a default policy is automatically assigned. A system may have multiple storage policies. For instance, all virtual machines that include PROD-SQL in their name or resource group might be set at RAID 1 and a five-percent read-cache reservation, and TEST-WEB would be automatically set to RAID 0.

### Dynamic policy changes

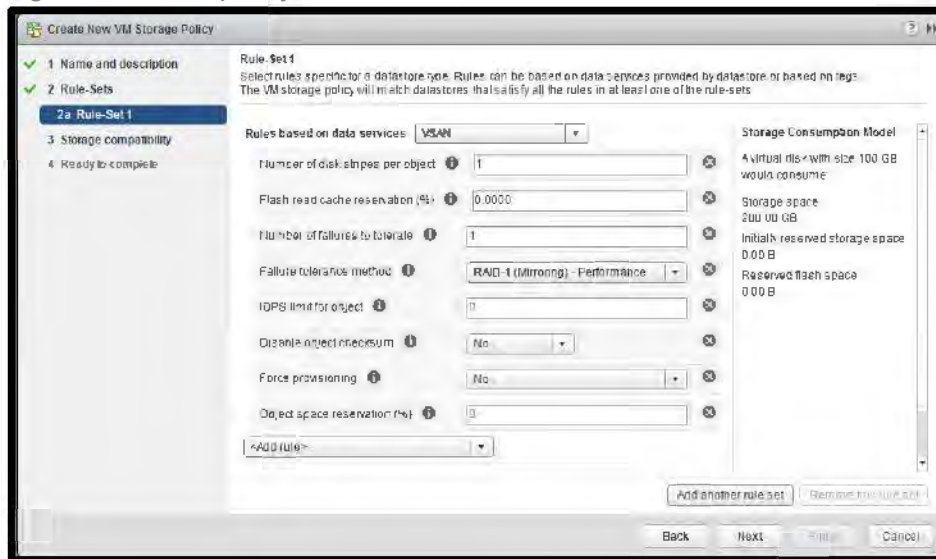
Administrators can dynamically change a VM storage policy. When changing attributes such as the number of Failures to tolerate, vSAN attempts to find a new placement for a replica with the new configuration. In some cases, existing parts of the current configuration can be reused, and the configuration just needs to be updated or extended. For example, if an object currently uses Failures to tolerate=1, and the user asks for Failures to tolerate=2, vSAN can simply add another mirror (and witness).

In other cases, such as changing the stripe width from one to two, vSAN cannot reuse existing replicas, and it creates a brand new replica (or replicas) without impacting the existing objects. A storage object has a status of *Compliant* if the attributes match the policy. When changing a policy online, the object may show as “*not compliant*” while vSAN reconfigures the object.

## Storage policy attributes

The figure below displays an example of the rule set for a storage policy.

Figure 47 vSAN policy attributes



### Number of disk stripes per object

This rule establishes the minimum number of capacity devices used for striping each virtual machine replica. A value higher than 1 might result in better performance, but it also results in higher resource consumption. The default value is the minimum, 1; the maximum value is 12. The stripe size is 1MB.

vSAN may decide that an object needs to be striped across multiple disks without any stripe-width policy requirement. The reason for this can vary, but typically it occurs when a VMDK is too large to fit on a single physical drive. However, when a particular stripe width is required, then it should not exceed the number of disks available to the cluster.

### Flash read cache reservation

Flash cache reservation refers to flash capacity reserved as read cache for the virtual machine object, and it applies to hybrid configurations only. By default, vSAN dynamically allocates read cache to storage objects based on demand. As a result, no need typically exists to change the default 0 value for this parameter.

In very specific cases, when a small increase in the read cache for a single VM can provide a significant change in performance, it is an option. It should be used with caution to avoid wasting resources or taking resources from other VMs.

The default value is 0 percent. Maximum value is 100 percent.

### Failures to tolerate

This FTT option generally defines the number of host and device failures that a virtual machine object can tolerate. For  $n$  failures tolerated,  $n+1$  copies of the VM object area created and  $2n+1$  hosts with storage are required.

The default value is 1. Maximum value is 3.



When erasure coding is enabled for a cluster (by setting FTM=Capacity), RAID 5 is applied if the number of Failures to tolerate is set to 1, and RAID 6 is applied if the number of Failures to tolerate is set to 2. Note a vSAN cluster requires a minimum of four nodes for RAID 5 and six nodes for RAID 6.

### Failure tolerance method

Failure tolerance method (FTM) specifies whether the data replication method optimizes for performance or capacity. The RAID 1 failure tolerance method provides better performance and consumes less memory and network resources but uses more disk space. RAID 5/6 erasure coding provides more usable capacity but consumes more CPU and network resources. (An upcoming section on erasure coding section provides additional information.)

### IOPS limit for object (QoS)

This attribute establishes Quality of Service (QoS) for an object by defining an upper limit on the number IOPS a VM/VMDK can perform. The default behavior is that all objects are not limited, and low priority applications could potentially consume resources in a manner that could impact more important workloads. This rule allows different limits for different applications and can be used to keep workloads from impacting each other (the noisy-neighbor issue) or to supply service-level agreements (SLAs) for different workloads and maintain performance for tier-1 applications.

IOPS is calculated based on all of the operations on this VM/VMDK (read & write) including snapshots.

A few notes regarding IOPS limits for objects:

When calculating IOPS, read and write are considered equivalent, but keep in mind that cache-hit ratio and sequentially are not considered.

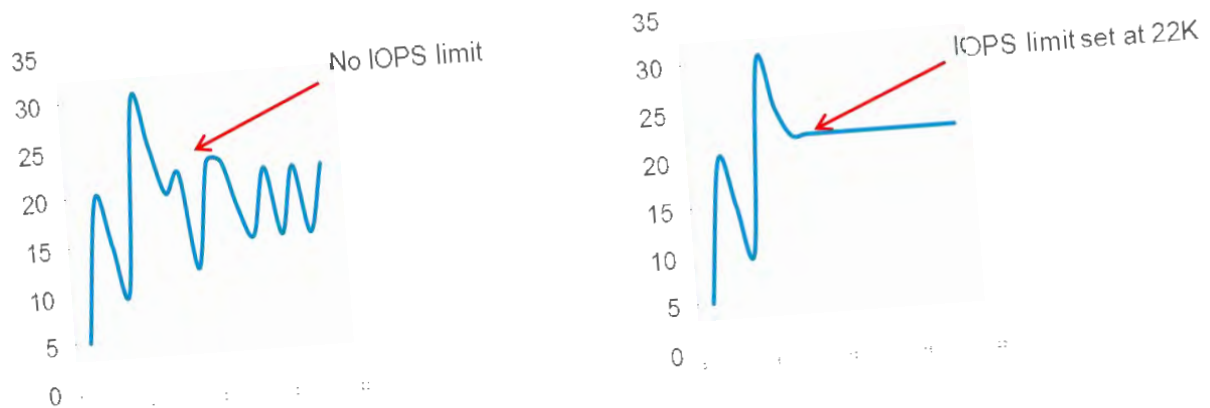
When an object exceeds its disk IOPS limit, I/O operations are throttled.

If the IOPS limit for object is set to 0, IOPS limits are not enforced.

vSAN allows the object to double the IOPS-limit rate during the first second of operation or after a period of inactivity.

The figure below illustrates how IOs are throttled when the IOPS limit policy is set.

Figure 48 IOPs limits establishes QoS



## Disable object checksum

vSAN uses end-to-end checksum to ensure data integrity by confirming that each copy of an object's 4K chunk of data is exactly the same. The checksum is five bytes in length and persists with the data. When data is written, the checksum is verified to ensure no corruption occurred when the data transverses the network. When data is read, checksum verifies that the data read matches what was written. If an error is detected; vSAN uses a replica to recover the data and logs an error. Checksum calculation and error-correction are transparent to the user.

The default setting for all objects in the cluster is On, which means that checksum is enabled. The best practice is not to change this setting because detecting data corruption is a critical and valuable feature of vSAN and should not be disabled.

## Force provisioning

If this option is set to Yes, the object is provisioned even if the rules specified in the storage policy cannot be satisfied by the datastore.

This parameter is sometimes used during an outage when resources are limited and normal provisioning policy cannot be satisfied.

The default is No and will not allow an object to be created if the policy rules cannot be satisfied. This is appropriate for most production environments. If set to Yes, an object can be created even not enough resources are available to satisfy the policy rules. In this case, the object will be displayed as Not Compliant.

## Object space reservation

Object space reservation is specified as a percentage of the total object size. It reflects the reserved, thick-provisioned space required for deploying virtual machines.

The default value is 0%. Maximum value is 100%.

The value should be set either to 0% or 100% when using RAID-5/6 in combination with deduplication and compression.

## Sparse Swap

Sparse Swap is a vSAN feature worth highlighting for its space efficiency, and it is available for both all-flash and hybrid environments. By default, vSAN swap files are thick provisioned—created with 100 percent space reservation. (For example, if the VM has 4GB of RAM, then the swap file also has 4GB of RAM.) While this guarantees sufficient capacity for the VM, it can consume too much memory—especially in large clusters with a lot of virtual machines. When Sparse Swap is enabled as an advanced host setting, the VM reserves less than 100 percent of memory space for the swap objects.

Sparse Swap is established by enabling the *Swap thick provision disabled* setting. Sparse swap files have an FTT setting of 1 and a FTM setting of RAID1 (Mirroring).

Leveraging Sparse Swap is an effective way to reduce memory overhead in environments that are not already over-committing memory, especially for all-flash VxRail environments where swap files represent a sizeable portion of the total required datastore capacity.



## I/O paths and caching algorithms<sup>4</sup>

This section elaborates on some of the vSAN concepts that have been introduced so far with additional, general information about vSAN caching algorithms. The next paragraphs briefly describe how vSAN leverages flash, memory, and rotating disks. They also illustrate the I/O paths between the guest OS and the persistent storage areas.

### Read caching

Each disk group contains an SSD drive used as a cache tier. On a hybrid system, 70 percent of the capacity is used by default for read cache (RC). The most active data is maintained in the cache tier and improves performance by minimizing the latency impact of reading from mechanical disk.

The RC is organized in terms of cache lines. They represent the unit of data management in RC, and the current size is 1MB. Data is fetched into the RC and evicted at cache-line granularity. In addition to the SSD read cache, vSAN also maintains a small in-memory (RAM) read cache that holds the most-recently accessed cache lines from the RC. The in-memory cache is dynamically sized based on the available memory in the system.

vSAN maintains in-memory metadata that tracks the state of the RC (both SSD and in memory), including the logical addresses of cache lines, valid and invalid regions in each cache line, aging information, etc. These data structures are designed to compress for efficiencies, using memory space without imposing a substantial CPU overhead on regular operations. No need exists to swap RC metadata in or out of persistent storage. (This is one area where VMware holds important IP.)

Read-cache contents are not tracked across power-cycle operations of the host. If power is lost and recovered, then the RC is re-populated (warmed) from scratch. So, essentially RC is used as a fast storage tier, and its persistence is not required across power cycles. The rationale behind this approach is to avoid any overheads on the common data path that would be required if the RC metadata was persisted every time RC was modified—such as cache-line fetching and eviction, or when write operations invalidate a sub-region of a cache line.


---

<sup>4</sup> Much of the content in this *specific* section has been extracted from an existing technical whitepaper: *An overview of VMware vSAN Caching Algorithms*.

## Anatomy of a hybrid read

Read operations follow a defined procedure. To illustrate, the VMDK in the example below has two replicas on esxi1 and esxi3. See the figure below.

Figure 49 Hybrid read



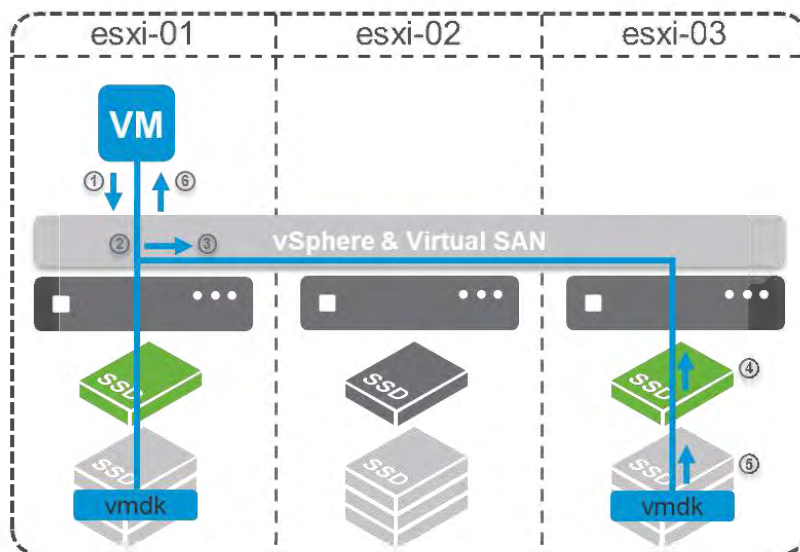
1. Guest OS issues a read on virtual disk
2. Owner chooses replica to read from
  - Load balance across replicas
  - Not necessarily local replica (if one)
  - A block always reads from same replica
3. At chosen replica (esxi-03): read data from flash write buffer, if present
4. At chosen replica (esxi-03): read data from flash read cache, if present
5. Otherwise, read from HDD and place data in flash read cache
  - Allocate a 1MB buffer for the missing cache line and replace “coldest” data (eviction of coldest data to make room for new read)
  - Each missing line is read from the HDD as multiples of 64KB chunks, starting with the chunks that contain the referenced data
6. Return data to owner
7. Complete read and return data to VM
8. Once the 1MB cache line is added to the in-line read cache, its population continues asynchronously. This occurs to explore both the spatial and temporal locality of reference, increasing the chances that the next reads will find in the read cache.

## Anatomy of an All-Flash read

1. Guest OS issues a read on virtual disk
2. Owner chooses replica to read from
  - Load balance across replicas
  - Not necessarily local replica (if one)
3. At chosen replica (esxi-03): read data from flash write buffer, if present
4. Otherwise, read from capacity flash device
5. Return data to owner
6. Complete read and return data to VM

The operation is shown in the figure below.

Figure 50 All-flash read



The major difference is that read-cache misses cause no serious performance degradation. Reads from flash capacity devices should be almost as quick as reads from the cache SSD. Another significant difference is that no need exists to move the block from the capacity layer to the cache layer, as in hybrid configurations.

## Write caching

In hybrid configurations write-back caching is done entirely for performance. The aggregate-storage workloads in virtualized infrastructures are almost always random, because of the statistical multiplexing of the many VMs and applications that share the infrastructure.

HDDs can perform only a small number of random I/O with a high latency compared to SSDs. So, sending the random write part of the workload directly to spinning disks can cause performance degradation. On the other hand, magnetic disks exhibit decent performance for sequential workloads. Modern HDDs may exhibit sequential-like behavior and performance even when the workload is not perfectly sequential. "Proximal I/O" suffices.

In hybrid disk groups, vSAN uses the write-buffer (WB) section of the SSD (by default, 30 percent of device capacity), as a write-back buffer that stages all the write operations. The key objective is to de-stage written data (not individual write operations) in a way that creates a benign, near-sequential (proximal) write workload for the HDDs that form the capacity tier.

In all-flash disk groups, vSAN utilizes the tier-1 SSD entirely as a write-back buffer (100 percent of the device capacity—up to a maximum of 600GB). The purpose of the WB is quite different in this case. It absorbs the highest rate of write operations in a high-endurance device and allows only a trickle of data to be written to the capacity flash. This approach allows low-endurance, larger-capacity SSDs for capacity.

Nevertheless, capacity SSDs are capable of serving very large numbers of read IOPS. Thus, no read caching occurs, except when the most-recent data referenced by a read operation still resides in the WB.

In both hybrid and All-flash, every write operation is handled through transactional processes: A record for the operation is persisted in the transaction log in the SSD.

The data (payload) of the operation is persisted in the WB.

Updated in-memory tables reflect the new data and its logical address space (for tracking) as well as its physical location in the capacity tier.

The write operation completes upstream after the transaction has committed successfully.

Under typical steady-state workloads, the log records of multiple write operations are coalesced before they are persisted in the log. This reduces the amount of metadata-write operations for the SSD. By definition, the log is a circular buffer, written and freed in a sequential fashion. Thus write amplification can be avoided (good for device endurance). The WB region allocates blocks in a round-robin fashion, keeping wear leveling in mind.

Even when a write operation overwrites existing WB data, vSAN never rewrites an existing SSD page in place. Instead, it allocates a new block and updates metadata to reflect that the old blocks are invalid. vSAN fills an entire SSD page before it moves to the next one. Eventually, entire pages are freed when all their data is invalid. (It is very rare to re-buffer data to allow SSD pages to be freed).

Also, because the device firmware does not have visibility into invalidated data, it sees no “holes” in pages. In effect, internal write leveling (by moving data around to fill holes in pages) is all but eliminated. This extends the overall endurance of a device. The vSAN design has gone to great lengths to minimize unnecessary writes to maximize cache SSD endurance. As a result, the life expectancy of SSDs implemented in vSAN may exceed the manufacturers’ specifications, which are developed with more generic workloads in mind.


### Anatomy of a write I/O—hybrid and All-Flash (FTM=mirroring)

1. VM running on host esxi-01.
2. esxi-01 is owner of virtual disk object.
  - *Failures to tolerate = 1*
3. Object has two (2) replicas on esxi-01 and esxi-03.
4. Guest OS issues write op to virtual disk.

5. Owner clones write operation.
  - In parallel: sends write op to esxi-01 (locally) and esxi-03
6. esxi-01, esxi-03 persist operation to flash (log).
7. esxi-01, esxi-03 ACK-write operation to owner.
8. Owner waits for ACK from both writes and completes I/O.
9. Later, backend hosts commit batch of writes.

The figure below illustrates the operation.

Figure 51 Hybrid and flash write I/O



### Distributed caching considerations


vSAN caching algorithms and data-locality techniques reflect a number of objectives and observations pertaining to distributed caching:

vSAN exploits temporal and spatial locality for caching.

vSAN implements a distributed, persistent cache on flash across the cluster. Caching is done in front of the disks where the data replicas live, not on the client side. A distributed-caching mechanism results in better overall flash-cache utilization.

Another benefit of distributed caching is during VM migrations, which can happen in some datacenters over ten times a day. With DRS and vMotion, VMs can move around from host to host in a cluster. Without a distributed cache, the migrations would have to move around a lot of data and rewarm caches every time a VM migrates. As the figure below illustrates, vSAN prevents any performance degradation after a VM migration.

Figure 52 vSAN prevents performance degradation after VM migration



The network introduces a small latency when accessing data on another host. Typical latencies in 10GbE networks range from 5 – 50 microseconds. Typical latencies of a flash drive, accessed through a SCSI layer, are near 1ms for small (4K) I/O blocks. So, for the majority of the I/O executed in the system, the network impact adds near 0.1 percent to the latency.

Few workloads are actually cache friendly, meaning that they don't take advantage of the way small increases in cache size can significantly increase the rate of I/O. These workloads can benefit from local cache, and enabling the Client Cache would be the right approach.

vSAN works with a View Accelerator (deduplicated, in-memory read cache), which is notably effective for VDI use cases.

vSAN features Client Cache that leverages DRAM memory local to the virtual machine to accelerate read performance. The amount of memory allocated is anywhere from .4 percent to 1GB per host.

## vSAN high availability and fault domains

vSAN policy attributes establish parameters to protect against node failures, but they may not be the most effective or efficient way to build tolerance for events like rack failures. This section reviews the availability features for vSAN clusters on the VxRail system. It starts out by looking at the availability implications on small VxRail deployments with fewer than four nodes.

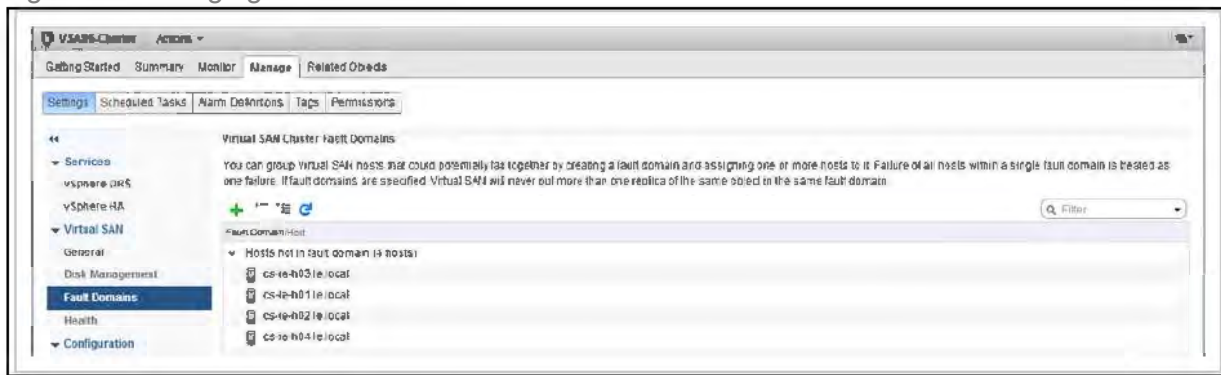
### Fault domain overview

vSAN and VxRail systems use fault domains as a way of configuring tolerance for rack and site failures. By default, a node is considered a fault domain. vSAN will spread components across fault domains, therefore, by default vSAN will spread components across nodes. Consider, for example, a cluster with four (4) four-node VxRail systems, each VxRail system placed in a different rack. By explicitly defining each four-node system as separate fault domains, vSAN will spread redundancy components across the different racks.

In terms of implementation, any host that is not part of another fault domain is considered its own single-host fault domain. VxRail requires at least three fault domains, and each has at least one host. Fault domain definitions recognize the physical hardware constructs that represent the domain itself. Once the domain is enabled, vSAN applies the active virtual machine storage policy to the entire domain, instead of just to the individual hosts. The number of fault domains in a cluster is calculated based on the FTT attribute: (Number of fault domains) = 2 \* (Failures to tolerate) + 1. Administrators can manage fault domains from the vSphere web client (as shown in the figure below.)



Figure 53 Managing fault domains



### Fault domains and rack-level failures

The fault domain mechanism detects when the configuration is vulnerable. Consider a cluster that contains four server racks, each with two nodes. If the FTT is set to 1, and fault domains are not enabled, vSAN might store both replicas of an object with hosts in the same rack. In that case, applications are exposed to a potential rack-level failure. With fault domains enabled, vSAN ensures that each protection component (replicas and witnesses) is placed in a separate fault domain, making sure that the nodes cannot fail together.

The figure below illustrates a four-rack setup, each with two ESXi nodes (a subset of the available hosts in a VxRail system). There are four defined fault domains:

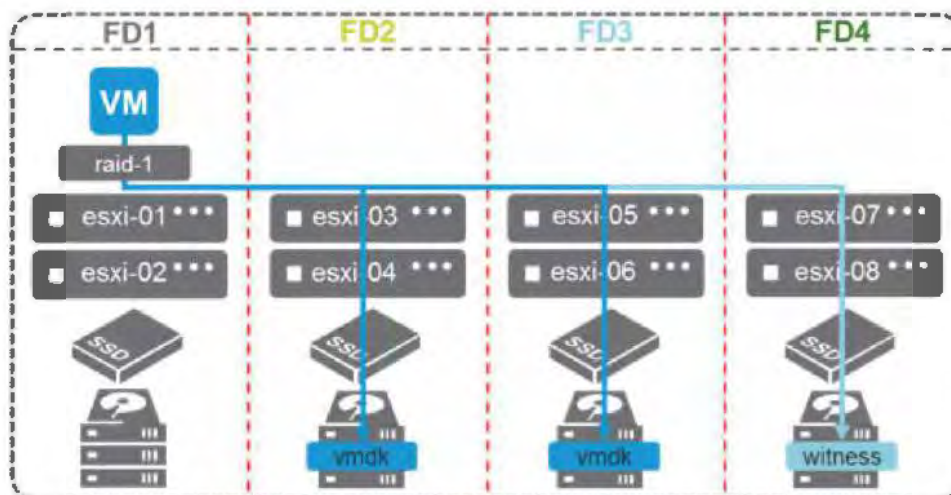
FD1 = esxi-01, esxi-02

FD2 = esxi-03, esxi-04

FD3 = esxi-05, esxi-06

FD4 = esxi-07, esxi-08

Figure 54 Fault domains for a four-rack VxRail configuration



This configuration guarantees that the replicas of an object are stored in hosts of different rack enclosures, ensuring availability and data protection in case of a rack-level failure.

## Cautions when deploying a minimum cluster configuration

When deploying a cluster that just meets the minimum requirements, it is important to understand the high availability implications. Choosing a 3-node minimum configuration for RAID-1 protection or a 4-node minimum configuration for RAID-5 protection means that a cluster will not be able self-heal by rebuilding data on another host if one host fails. When a host is in maintenance mode such as a node upgrade, the data is exposed to a potential failure or inaccessibility if an additional failure occurs.

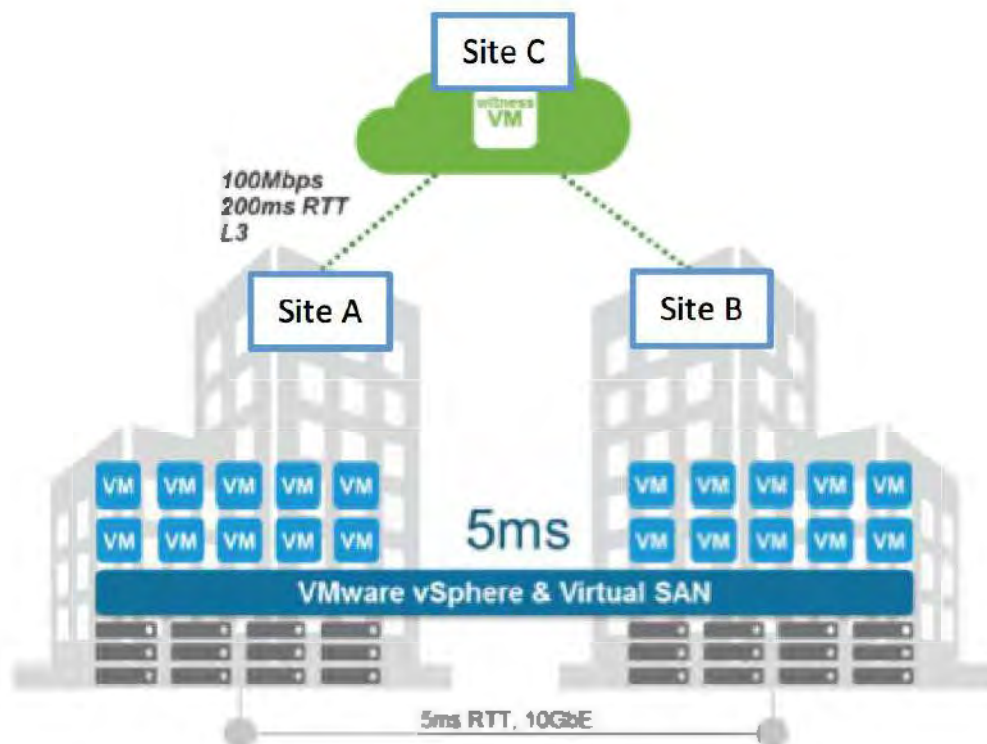
## vSAN Stretched Cluster

The concept of a stretched cluster is a good example of vSAN's native integration with vSphere. With VxRail, stretch clustering extends availability of large enterprise datacenter. The stretched cluster is a specific configuration implemented in environments where the requirement for datacenter-level downtime avoidance is absolute. Similar to how fault domains enable “rack awareness” for rack failures; stretched clusters provide “datacenter awareness,” maintaining virtual machine availability despite specific datacenter failure scenarios.

In a VxRail environment, stretched clusters with a witness host refers to a deployment where a vSAN cluster consists of two active/active sites with an identical number of ESXi hosts distributed evenly between them. The sites are connected via a high bandwidth/low latency networking.

In the figure below, each site is configured as a vSAN fault domain. The nomenclature used to describe the stretched cluster configuration is X+Y+Z, where X is the number of ESXi hosts at Site A, Y is the number of ESXi hosts at Site B, and Z is the number of witness hosts at site C.

Figure 55 Stretched VxRail cluster



A virtual machine deployed on a stretched cluster has one copy of its data on Site A, and another on Site B, as well as witness components placed on the host at Site C.



It is a singular configuration, achieved through a combination of fault domains, hosts and VM groups, and affinity rules. In the event of a complete site failure, the other site still has a full copy of virtual machine data and at least half of the resource components are available. That means all the VMs remain active and available on the vSAN datastore. The recovery point objective (RPO) is zero and the data recovery time objective (RTO) is zero. The application RTO is dependent upon the application recoverability.

The minimum configuration supported by VxRail is 3+3+1 (7 nodes); the maximum is 15+15+1 (31 nodes). Stretched clusters are supported by both hybrid and all-flash VxRail configurations. Stretched clusters running in the latest VxRail software versions support customer-driven upgrades. For clusters running older VxRail software, customers will need to contact support to facilitate upgrades.

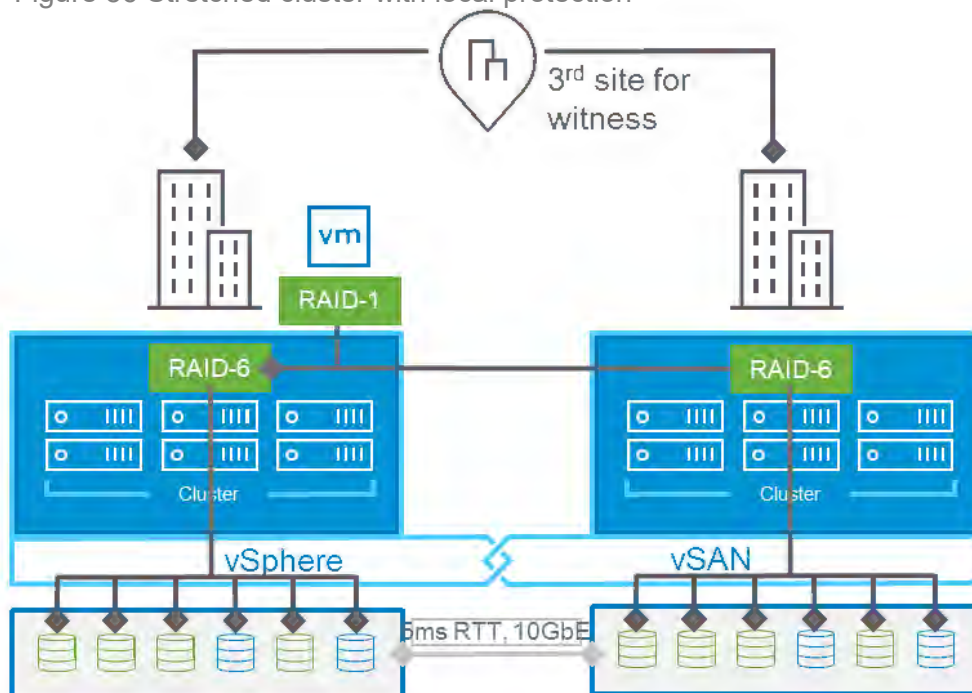
For more information, refer to *VxRail vSAN Stretched Clusters Planning Guide*:  
<https://vxrail.is/stretchedclusterplanning>

### Stretched Cluster with Local Protection

VxRail software version 4.5 and vSAN 6.6 or above support Stretched Clusters with Local Protection. This feature mirrors data between sites, with each site applying local data protection for increased protection. The protection is specified using two parameters: Primary Failures To Tolerate (PFTT), and Secondary Failures To Tolerate (SFTT). PFTT refers to the protection between sites which is always RAID1 mirroring. SFTT is the local protection applied at each site. Hybrid configurations support SFTT of 0, 1, 2, or 3 with RAID1 (mirroring) Failure Tolerance Method (FTM). All-flash configurations support SFTT of 0, 1, 2, or 3 with RAID1 FTM or SFTT of 1 or 2 with Erasure Coding FTM.

Local protection for an all-flash stretched cluster configuration is shown in the figure below.

Figure 56 Stretched cluster with local protection



## Site locality

In a conventional storage-cluster configuration, reads are distributed across replicas. In a stretched cluster configuration, the vSAN Distributed Object Manager (DOM) also takes into account the object's fault domain, and only reads from replicas in the same domain. That way, it avoids any lag time associated with using the inter-site network to perform reads.

## Networking

Both Layer-2 (same subnet) and Layer-3 (routed) configurations are used for stretched cluster deployments. A Layer-2 or Layer-3 connection is configured between data sites, and Layer-3 connection between the witness and the data sites.

The bandwidth between data sites depends on workloads, but Dell EMC requires a minimum of 10Gbps for VxRail systems in a stretched cluster configuration. The supported latency for witness hosts is up to 200ms RTT and a bandwidth of 2Mbps for every 1,000 vSAN objects. Also bear in mind that the latency between data sites should be no be greater than 5ms, generally estimated at 500km or about 310 miles.

## Stretched cluster heartbeats and site bias

Stretched cluster configurations effectively have three fault domains. The first functions as the preferred data site, the second is the secondary data site, and the third is simply the witness host site.

The vSAN master node is placed on the preferred site and the vSAN backup node is placed on the secondary site. As long as nodes (ESXi hosts) are available in the preferred site, then a master is always selected from one of the nodes on this site—similarly for the secondary site, as long as nodes are available on the secondary site.

The master node and the backup node send heartbeats every second. If heartbeat communication is lost for five consecutive heartbeats (five seconds), the witness is deemed to have failed. If the witness has suffered a permanent failure, a new witness host can be configured and added to the cluster. Preferred sites gain ownership in case of a partition.

After a complete failure, both the master and the backup end up at the sole remaining live site. Once the failed site returns, it continues with its designated role as preferred or secondary, and the master and secondary migrate to their respective locations.

In terms of the communication with the witness, if the heartbeat pauses for five consecutive beats, the master assumes that the witness failed. If it's a permanent failure, a new witness host needs to be configured and added to the cluster, and preferred sites gain ownership in case of a partition.

## vSphere HA settings for stretched cluster

A stretched cluster requires the following vSphere HA settings:

Host monitoring is enabled by default in all VxRail deployments, including of course stretched cluster configurations. This feature also uses network heartbeat to determine the status of hosts participating in the cluster. It indicates a possible need for remediation, such as restarting virtual machines on other cluster nodes.

Configuring admission control ensures that vSphere HA has sufficient available resources to restart virtual machines after a failure. This may be even more significant in a stretched cluster than it is in a single-site cluster, because it makes the entire, multi-site infrastructure resilient.

Workload availability is perhaps the primary motivation behind most stretched cluster implementations.

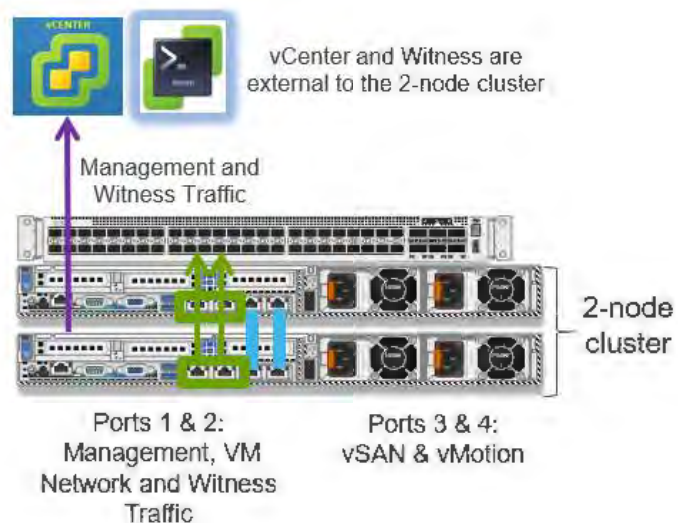
The deployment needs sufficient capacity to accommodate a full site failure. Since the stretched cluster equally divides the number of ESXi hosts between sites, Dell EMC recommends configuring the admission-control policy to 50 percent for both CPU and memory to ensure that all workloads can be restarted by vSphere HA.

## 2-Node Configuration

VxRail software version 4.7.100 introduces support for a fixed 2-node configuration that uses direct connection between the two nodes for cluster operations. With its small footprint, it can be an appropriate, cost-effective solution for locations with limited space and workload requirements. The support requires that the configuration to be a brand-new deployment which means existing clusters cannot utilize node removal to convert into a 2-node configuration. Cluster expansion is not supported therefore this solution should be targeted for specific use cases. Only the VxRail E560 and E560F are supported in a 2-node cluster. Users can still benefit from automated lifecycle management.

The configuration consists of the two nodes directly connected and a Witness to provide quorum for the cluster. The direct connection is for the vSAN and vMotion traffic. The Witness is a virtual appliance installed on an ESXi host which must reside outside of the 2-node cluster, i.e. in another datacenter or a physical host in the same rack/location. The Witness has individual connections to both nodes which requires VLANs to separate Witness management traffic from vSAN traffic. The configuration only supports mirroring (FTT=1). Witness host is used as the tiebreaker. Each node and the Witness is an individual fault domain for a total three in the cluster.

Figure 57 VxRail 2-node cluster



A special workflow in the First Run experience is used to deploy the 2-node cluster. The workflow includes the setup of the Witness appliance and Witness traffic separation. The configuration must use a customer-supplied vCenter for management. With only two data nodes in the cluster, users need be cognizant of the cluster load to prevent data unavailability in case of a node failure and a single node servicing the entire cluster workload.

The 2-node cluster supports per-socket vSAN Standard, Advanced, and Enterprise license editions. Refer to the [vSAN Licensing Guide](#) for more details.

For more information about this configuration, refer to the [VxRail Technical Deck](#).

## Snapshots

Snapshots have been around for a while as a means of capturing the state of a data object at a particular point-in-time (PIT), so that it can be rolled back to that state if needed after a logical or physical failure. In the case of the VxRail solution, administrators can create, roll back, or delete VM snapshots using the Snapshot Manager in the vSphere web client. Each VM supports a chain of up to 32 snapshots.

A virtual machine snapshot generally includes the settings (.nvram and .vmx), the power state of all the VM's associated disks, and optionally, the memory state. Specifically, each snapshot includes:

**Delta disk:** The state of the virtual disk at the time the snapshot is taken is preserved. When this occurs, the guest OS is unable to write to its .vmdk file. Instead, changes are captured in an alternate file named VM\_name-delta.Vmdk.

**Memory-state file:** VM\_name-Snapshot#.Vlms, where # is the next number in the sequence, starting with 1. This file holds the memory state since the snapshot was taken. If memory is captured, the size of this file is the size of the virtual machine's maximum memory. If memory is not captured, the file is much smaller.

**Disk-descriptor file:** VM\_name-00000#.vmdk, a small text file that contains information about the snapshot.

**Snapshot-delta file:** VM\_name-00000#-delta.Vmdk, which contains the changes to the virtual disk's data at the time the snapshot was taken.

**VM\_name.Vmsd:** This snapshot list file is created when virtual machine itself is deployed. It maintains VM snapshot information that goes into a snapshot list in the vSphere web client. This information includes the name of the snapshot .Vlms file and the name of the virtual-disk file.

The snapshot state uses a .Vlms extension and stores the requisite VM information at the time of the snapshot. Each new VM snapshot generates a new .vmsn file. The size of this file varies, based on the options selected during creation. For example, including the memory state of the virtual machine increases the size of the .vmsn file. It typically contains the name of the VMDK, the display name and description, and an identifier for each snapshot.

Other files might also exist. For example, a snapshot of a powered-on virtual machine has an associated snapshot\_name\_number.vmem file that contains the main memory of the guest OS, saved as part of the snapshot.

A quiesce option is available to maintain consistent point-in-time copies for powered-on VMs. VMware tools may use their own sync driver or use Microsoft's Volume Shadow Copy Service (VSS) to quiesce not only the guest OS files system, but also any Microsoft applications that understand VSS directives.

## Storage efficiency using deduplication and compression

Storage capacity requirements continue to grow exponentially, and IT organizations are looking for ways to increase storage efficiency in order to meet their growing capacity requirements at the lowest cost. One way to do this is to use data deduplication and compression, which can result in more capacity at a lower cost. Many environments can achieve an effective capacity that is twice the raw capacity. These data reduction capabilities can be utilized on all-flash VxRail clusters.

Compression and deduplication techniques have been in place for a number of years but have not been widely adopted because of the overhead and system resources required to implement. Today, VxRail all-flash models with many cores and lots of memory per processor are a

powerhouse! Along with the architectural efficiencies of vSAN, the space savings more than offset the slight overhead. A VxRail all-flash configuration often provides more effective capacity at a lower cost than a hybrid HDD solution.

With vSAN, deduplication and compression occurs inline when data is de-staged from the cache to the capacity drives. First data is deduplicated by removing redundant copies of blocks that contain the exact same data. This is done at the 4K block level.

The figure below shows a typical virtual machine environment.

Figure 58 Inline data deduplication



While all VMs are unique, they share some amount of common data. Rather than saving multiple copies of the same data, identical blocks are saved once on SSD, and references to the unique blocks are tracked using metadata that is maintained in the capacity tier.

The deduplication algorithm is applied at the disk-group level and results in only a single copy of each unique 4K block per disk group. While duplicated data blocks may exist across multiple disk groups, by limiting the deduplication domain to a disk group, a global lookup table is not required, which minimizes network overhead and CPU utilization.

LZ4 compression is applied after the blocks are deduplicated and before being written to SSD. If the compression results in a block size of 2KB or less, the compressed version of the block is persistently saved on SSD. If the compression does not result in a block size of less than 2KB, the full 4K block is written to SSD.

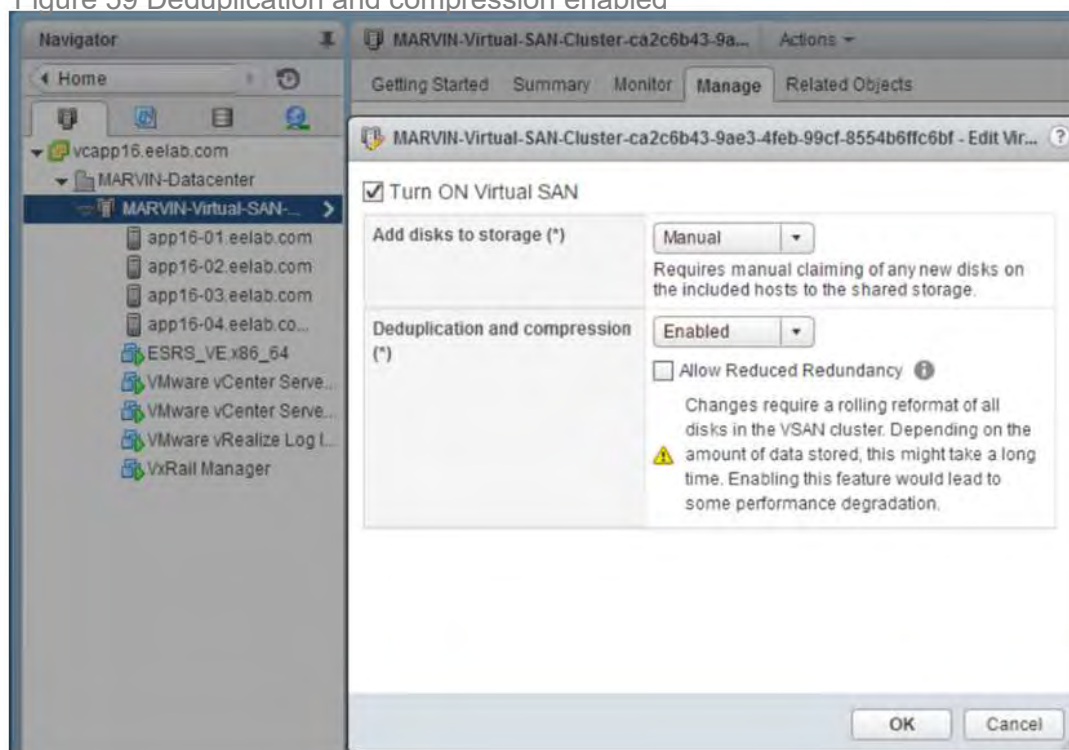
Almost all workloads benefit some from deduplication. However typical virtual server workloads with highly redundant data such as full clone virtual desktops or homogenous server operating systems benefit most. Compression provides further data reduction. Text, bitmap, and program files are very compressible, and 2:1 is often possible. Other data types that are already compressed, such as certain graphics formats and video files or encrypted files, may yield little or no reduction.

Deduplication and compression are disabled by default and are enabled together at the cluster level. (See the figure below.) While it can be enabled at any time, enabling it when the system is



initially setup is recommended to avoid the overhead and potential performance impact of having to deduplicate and compress existing data through post processing rather than to do it inline.

Figure 59 Deduplication and compression enabled



### Deduplication and compression overhead

Deduplication algorithms break data files into contiguous segments, or compute fingerprints, used to identify duplicate segments and reduce the data footprint. This is a basic deduplication concept. The specific approach varies among system vendors, but any deduplication method consumes CPU to compute the segment fingerprints or hash keys, and it executes I/O operations when performing lookups on the segment index tables.

vSAN computes the fingerprints and looks for duplicated segments only when the data is being de-staged from the cache to the capacity tier. This means that under normal operations, VM writes to the write buffer in the cache SSD should not incur any latency impact.

The cost of the deduplication occurs when data is de-staged from the cache to the capacity tier. It consumes a portion of CPU capabilities reserved for vSAN, and the disk operations generated by the index lookups consumes a portion of the backend I/O capabilities.

Because resource consumption varies according to I/O patterns, data types and so on, consult with an Dell EMC or VMware specialist before deciding whether deduplication is recommended for your application.

More information can be found in *Technical Whitepaper VMware VSAN 6.2 Space Efficiency Technologies* at <http://www.vmware.com/files/pdf/products/vsan/vmware-vsan-62-space-efficiency-technologies.pdf>.

## Erasure coding

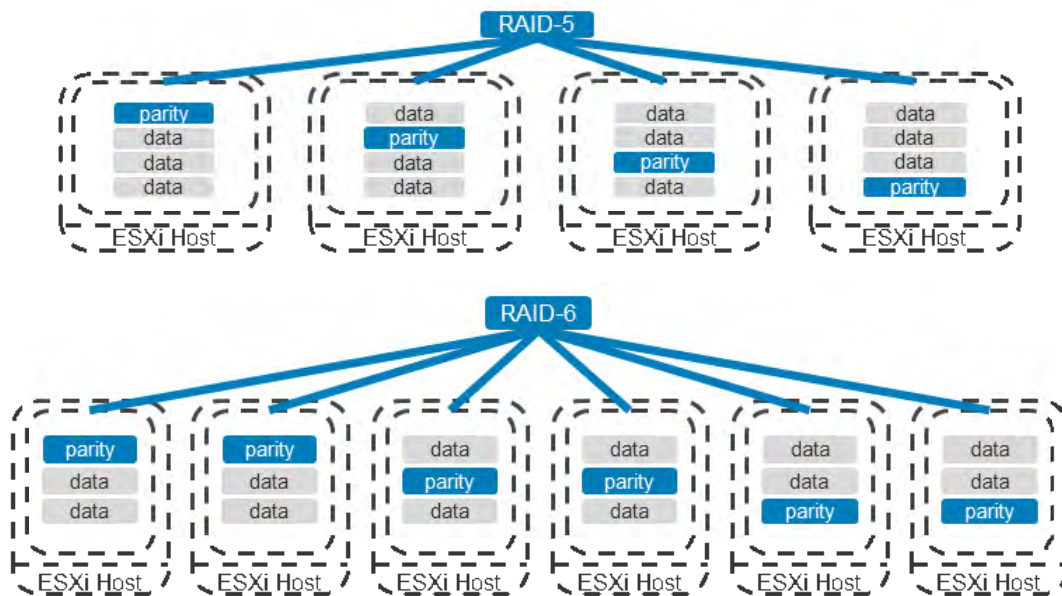
Erasure coding is another capacity-efficient solution for the failure tolerance method and data protection on all-flash VxRail configurations. As an alternative failure tolerance method to the data replication provided by RAID-1 mirroring, erasure codes can provide up to 50 percent more usable capacity than purely conventional RAID-1 mirroring, which drains storage space.

Erasure coding breaks up data into fragments and distributes redundant chunks of data across the system. It introduces redundancy by using data blocks and striping. To explain basically, data blocks are grouped in sets of  $n$ , and for each set of  $n$  data blocks, a set of  $p$  parity blocks exists. Together, these sets of  $(n + p)$  blocks make up a stripe. The crux is that any of the  $n$  blocks in the  $(n + p)$  stripe is enough to recover the entire data on the stripe.

In VxRail clusters, the data and parity blocks that belong to a single stripe are placed in different ESXi hosts in a cluster, providing a layer of failure tolerance for each stripe. Stripes don't follow a one-to-one distribution model. It is not a situation where the set of  $n$  data blocks sits on one host, and the parity set sits on another. Rather, the algorithm distributes individual blocks from the parity set among the ESXi hosts.

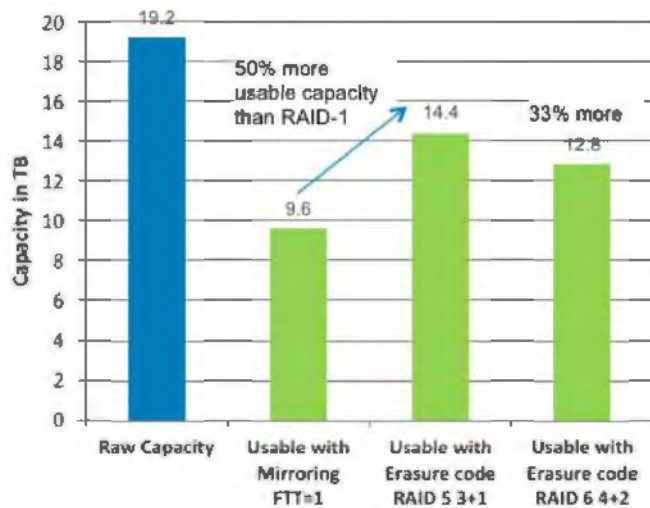
Erasure coding provides single-parity data protection (RAID-5) that can tolerate one failure (FTT=1) and double-parity data protection (RAID-6) that can tolerate two failures (FTT=2). The diagrams below illustrate the implementations. A single-parity stripe uses three data blocks and one parity block (3+1), and it requires a minimum of four hosts or four fault domains to ensure availability in case one of the hosts or disks fails (as shown below). It represents a 30 percent storage savings over RAID-1 mirroring. Dual parity saves as much as 50 percent capacity over RAID-1. It uses four data blocks plus two parity blocks (4+2) and requires a minimum of six nodes. See the figure below.

Figure 60 RAID-5 (FTT=1) requires a minimum of four nodes and RAID-6 (FTT=2) with 4+2 nodes



Look at the comparison of usable capacity in the figure below. The erasure-code protection method increases the usable capacity up to 50 percent compared to mirroring.

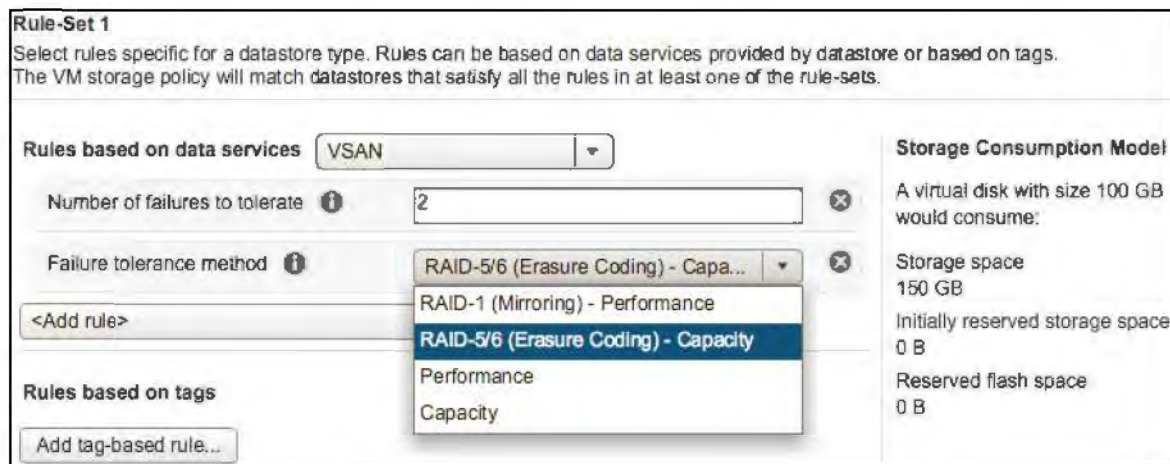
Figure 61 Erasure coding increases usable capacity up to 50 percent



## Enabling Erasure Coding

The SPBM policy Failure Tolerance Method (FTM) lets administrators choose between RAID-1 (Mirroring) and RAID-5/6 (Erasure Coding). The FTT policy determines the number of parity blocks written by the erasure code. See the figure below.

Figure 62 FTT policy determines the number of parity blocks written by the erasure code



VxRail implements erasure coding at a very granular level, and it can be applied to VMDKs, making for a nuanced approach. Configurations for VMs with write-intensive workloads—a database log, for instance—can include a mirroring policy, while the data component can include an erasure coding.

## Erasure coding overhead

Erasure coding saves space but increases backend overhead. Computing parity blocks consumes CPU cycles and adds overhead to the network and disks, as does distributing data slices across multiple hosts. This extra activity can affect latency and overall IOPS throughput.

The rebuild operation also adds overhead. In general, rebuild operations multiply the number of reads and network transfers used for replication. A formula is available here, too. If  $n$  refers to the number of blocks in a stripe, then the rebuild operations cost  $n$  times that of ordinary



replication. For a 3+1 stripe, that means three disk reads and three network transfers for every one of conventional data-replication. The rebuild operation can also be invoked to serve read requests for currently available data.

This additional I/O is the primary reason why only all-flash VxRail configurations use erasure coding. The rationale here is that the flash disks compensate for the extra I/O.

## vSAN Encryption

In vSAN 6.6, encryption is a new datastore level setting. vSAN Encryption is a datastore, or cluster-wide, level setting applied to all VM components in the cluster. This feature solves the concern of media theft. An advantage of the cluster wide setting is that deduplication and compression are applied prior to the encryption. This provides space savings benefit over the vSphere 6.5 Encryption option. Like vSphere Encryption, a KMIP-compliant Key Management Server like CloudLink or Hytrust must be used in conjunction with vSAN Encryption. vSAN encryption is FIPS 140-2 Level 1, AES 256 compliant.

## VxRail integrated software

VxRail has integrates software in two ways, both of which are fully engineered, tested, validated, manufactured and supported as a single offering from Dell EMC.

Products that are native to vSphere, including vSphere Replication

Dell EMC software products, including RecoverPoint for Virtual Machines

### VM Replication

Several options are available for replicating virtual machines for data protection and disaster recovery in VxRail clusters. Among them are solutions integrated into the VMware software stack: VMware vSphere Replication (VR), and RecoverPoint for Virtual Machines (RP4VM), which is built on enterprise proven Dell EMC RecoverPoint technology.

When choosing the right solution, Recovery Point Objectives (RPO) is an important consideration. RPO defines the maximum acceptable age of data recovered from a replicated copy as a result of a data loss issue. For example, if a virtual machine is deleted and the RPO for the virtual machine is 24 hours, a recovered copy of the virtual machine should contain all data except for any changes that occurred in the last 24 hours. Depending on the technology, RPOs as low as a few seconds or minutes can be configured.

Another consideration is the number of recovery points to retain. When configuring replication for a virtual machine, an administrator has the option to enable the retention of multiple recovery points (point-in-time instances). This can be useful when an issue is discovered several hours, or even a few days, after it occurred. For example, a replicated virtual machine with a four-hour RPO contracts a virus, but the virus is not discovered until six hours after infestation. As a result, the virus has been replicated to the target location. With multiple recovery points, the virtual machine can be recovered and then reverted to a recovery point retained before the virus issue occurred.

### VMware vSphere Replication

VMware vSphere Replication is a hypervisor-based, asynchronous replication solution that is fully integrated with VMware vCenter Server. It uses a proprietary replication engine developed by VMware and is included with VMware vSphere Essentials Plus Kit and higher editions of VMware vSphere. While VR works well with vSAN storage in a VxRail hyperconverged environment, it is completely independent of the underlying storage and allows for replication between heterogeneous storage types. This is useful when a VxRail system is part of a larger virtualization environment that includes SAN or other storage types. VR provides data protection locally within the same vCenter environment as well as disaster recovery and avoidance to another vCenter site. It also supports replication to Service Provider clouds such as vCloud Air.

The figure below shows VR replication between two VxRail sites.

Figure 63 VMware vSphere Replication with VxRail



In this example, several VMs are replicated to a remote site with a RPO of 15 minutes. The remote site maintains multiple VM images, allowing roll back to different points-in-time.

The VR components that transmit replicated data are built into vSphere and use Secure Sockets Layer (SSL) connection. A best practice is to isolate network traffic to improve performance and security. VR also includes one or more prebuilt Linux-based vSphere Replication vApps. The first virtual appliance is referred to as the vSphere Replication Management server. It receives replicated data, manages authentication, and maintains mappings between the source virtual machines and the replicas at the target location. Each appliance requires 18GB of VMDK storage, 4GB of memory, and either two or four virtual CPUs.

VR configuration includes specifying the Recovery Point Objectives (RPO) within a range of 15 minutes to 24 hours as well as the number of point-in-time time images to maintain. This is all done from within the vSphere web client using a simple wizard. Once replication has been configured for a virtual machine, vSphere Replication begins the initial full synchronization of the source virtual machine to the target location. The time required to complete this initial synchronization can vary depending on the amount of data that must be replicated and the amount of available network bandwidth. After the initial full synchronization, only changed data is transmitted in each replication cycle, minimizing the network bandwidth requirements.

The replicated data is first written to a redo log. After all changes for the current replication cycle have been received and written to the redo log, the data in the redo log is merged into the base image. This process ensures a consistent and recoverable VM image is available at all times.

vSphere Replication delivers flexible, reliable and cost-efficient replication for data protection and disaster recovery. It is seamlessly integrated within the VMware product stack for simple deployment, configuration, and management.

## Dell EMC RecoverPoint for Virtual Machines

RecoverPoint for Virtual Machines (RP4VM) is based on the RecoverPoint continuous data protection technologies that have been proven in enterprise environments for over ten years with over 350 million run hours and an entire exabyte of data protected. Like VR, RecoverPoint for VMs enables customers to replicate virtual machines simply and is configured and managed from within vSphere web client using the RecoverPoint plug-in.

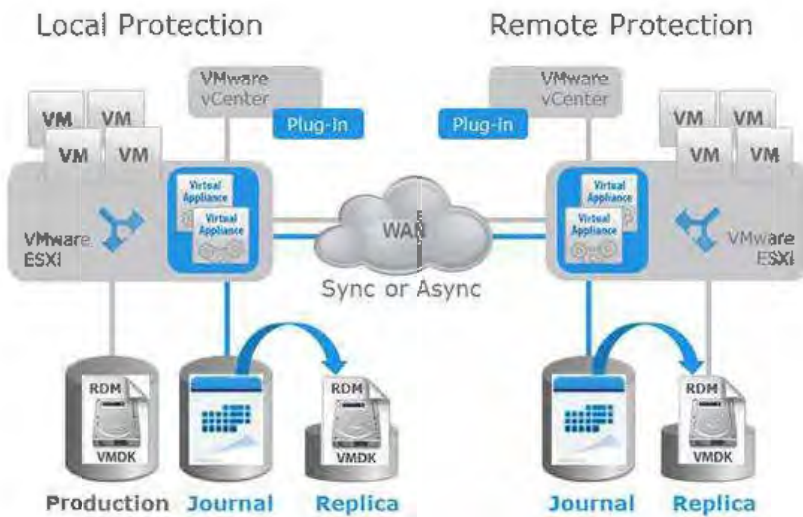
While similar in concept to VR, RP4VM has a number of unique capabilities including the ability to configure an RPO as low as zero seconds, as well as concurrent local and remote replication, and multi-site 2:1 and 1:2 configurations.

RecoverPoint for VM includes a RecoverPoint Virtual Appliance (vRPA) that manages all aspects of data replication. Two-to-eight vRPAs are grouped as a cluster. For local replication, a single cluster would be configured. For remote replication, a cluster is configured at both the local and remote sites. A RecoverPoint splitter is also installed on each ESXi server. The splitter takes each write and sends it to both the VMDK and to the vRPA. Management of RP4VM is all done from within vCenter using the RecoverPoint for VMs plugin.

The source of replication is the virtual machine, and associated application data and is referred to as the production copy. RP4VM performs continuous data replication of writes to the production copy. Each write is split and sent to both the VMDK and to a journal. The journal provides DVR-like roll back and is used to create a local copy that reflects any point in time. The local copy can be used to recover from logical errors and data corruption. Optionally, RP4VM can be configured to synchronously or asynchronously replicate to a remote vRPA cluster. The local and remote vRPA clusters communicate over a WAN connection and compression and de-duplication provides optimization and reduces bandwidth consumption.

The following figure shows an environment with both local and remote data protection and how writes are split, journaled, and used to create point-in-time images.

Figure 64 Continuous local and remote protection



Other RP4VM features include automated discovery, provisioning and orchestration of DR workflows, including test, and failover and failback of a single VM or multiple VMs using consistency groups.

A license Starter Pack is included with each VxRail system. Additional licenses can be purchased as needed.

## VxRail replication use case

Both VR and RPVM provide data protection and disaster-recovery capabilities. Both are software-only solutions that are embedded in the hypervisor, use vApps for control, and are managed from vCenter. Which one to use depends on the use case.

Local replication to other systems, remote replication to other VxRail clusters, or non-VxRail vSphere clusters may be configured. Remote replication may be used to provide DR to ROBO sites. Other use cases include VM migration from other vCenter environments into and between VxRail environments.

For basic replication that is fully integrated and easy to manage using VMware vCenter, VR provides these capabilities at no additional cost. If an RPO of less than five minutes, only RP4VM can provide continuous data replication with DVR-like playback. RPVM also supports more flexible deployment options including concurrent local and remote replication and multiple sites in a 2:1 or 1:2 configuration.

Regardless of the replication choice, solution sizing should include the additional storage capacity required and the overhead of replication vApps.

## Support for external network storage

The vSAN presents a robust, secure, and efficient shared datastore to all nodes within a VxRail cluster. External SAN storage is typically not part of a VxRail environment. However, often a requirement exists to access external storage in order to move virtual machines and data into a VxRail environment or move data between environments. Fibre Channel SAN connectivity is supported, and so is IP-based storage. . An important distinction is that data in the Fibre Channel, iSCSI or NFS datastore is self-contained and is not distributed to the disk groups within the VxRail cluster.

### Fibre Channel with VxRail

Customers can order Fibre Channel (FC) host bus adaptors (HBA) with their VxRail for external storage. FC storage can be configured to complement local VxRail cluster storage. Common use cases for external storage are customer's desire to continue to use their existing storage array as a secondary storage to VxRail, or they are looking for a method to migrate data from their FC storage to VxRail vSAN datastores. VxRail does not provide lifecycle management to the FC HBA. Customers will need to manage it via vCenter Server.

Using a FC HBA, customer can connect to storage arrays that is supported by the HBA card and validated by VMware. However, Dell EMC will only provide support for connection of the HBA to a Dell EMC storage array (i.e. Dell Compellent, Clariion/VNX/Unity, Symmetrix DMX, Symmetrix VMAX/PowerMAX, and XtremIO) that is qualified by eLab.

When configuring external storage via the FC HBA, customer is allowed to install VM/VIB/drivers to operationalize the use of the external storage as required. The customer is responsible for maintaining and updating it. Customers can install multiple HBAs if there are slots available on the PCIe bus.

### iSCSI with VxRail

iSCSI can be used to provide mobility for VMs and associated data onto and between VxRail environments. The figure below shows a VxRail environment that includes iSCSI storage in addition to the vSAN datastore.

Figure 65 Data mobility into and between VxRail environments



Data on the iSCSI storage is easily moved into the VxRail vSAN environment or between VxRail environments.

Existing iSCSI storage can also be used to provide additional capacity to the VxRail environment. However with the VxRail scale-up and scale-out configuration flexibility, external storage is typically not used to meet capacity requirements.

iSCSI provides block-level storage using the SCSI protocol over an IP network. SCSI uses a client-server, initiator-target model where initiators issue read/write operations to target devices, and targets either return the requested read data or persistently save write data. iSCSI in a VMware environment is standard functionality. A software adapter using the NIC on an ESXi host is configured as an initiator, and targets on an external storage system present LUNs to the initiators. The external LUNs could be used by ESXi as raw device mapping (RDM) devices, however usually, the use case is for VxRail to configure them as VMFS datastores. (Refer to vSphere documentation for more information: *Using ESXi with iSCSI SAN.*)


iSCSI configuration is performed using the vSphere web client. The steps involve creating a port group on the VDS, creating a VMkernel Network Adapter and associating it with the port group, and assigning an IP address. Then, from the vCenter Manage Storage Adapters view, the Add iSCSI Software Adapter dialog is used to create the software adapter. The last step is to bind the iSCSI software adapter with VMkernel adapter. Once this is complete, iSCSI targets and LUNs can be discovered and used to create new datastores and map them to the hosts in the cluster. (Refer to the VMware documentation for more details.)

iSCSI works best in a network environment that provides consistent and predictable performance, and a separate VLAN is usually implemented. iSCSI network requirements should be considered when planning the network requirements for VxRail environment to make sure connectivity to the external iSCSI storage system exists, and the additional network traffic will not impact other applications.

## NFS with VxRail

NFS is a network filesystem that provides file-level storage using the NFS protocol over an IP network. It can work in use cases similar to iSCSI—the difference being that NFS devices are presented as file systems rather than block devices. The figure below shows an NFS file system that has been exported from a network-attached server and mounted by the ESXi nodes in the VxRail environment.

Figure 66 Network-attached file system with VxRail



This enables data mobility into and between VxRail environments as well as enabling additional storage capacity.

The external NFS server can be an open system host, typically Unix or Linux, or a specially built system. The NFS server takes physical storage and creates a file system. The file system is exported and client systems, in this example ESXi hosts in a VxRail system, mount the file system and access it over the IP network.

Similar to iSCSI, NFS is a standard vSphere feature and is configured using the vCenter web client. This is done in the Hosts and Clusters view under Related Objects and the New Datastore dialog. Select NFS as datastore type, the NFS version, the name of the datastore, the



IP address or hostname of the NFS server that exported the filesystem, and the host that will mount it. The NFS filesystem will appear like the vSAN datastore. VMs, templates, OVA files, and other storage objects can be easily moved between the NFS filesystem and the vSAN datastore using vMotion.

As with iSCSI, NFS works best in network environments that provide consistent and predictable performance. The network requirements for NFS should be considered when initially planning the network requirements for VxRail environment.



# VxRail solutions and ecosystem

Dell EMC offers a full range of flexible consumption models that make it faster and easier for businesses to use VxRail to fuel digital transformation. These consumption models include both the technology itself and how businesses pay for this technology.

## VMware Validated Design with VxRail

VMware Validated Design (VVD) is a family of solutions for datacenter designs that span compute, storage, networking, and management, serving as a blueprint for your Software-Defined Datacenter (SDDC) implementation. The VVD is a reference architecture for how to deploy, operate, and maintain a VMware SDDC. VxRail is supported starting with the VVD 4.2.

The VMware VVD provides a framework for complete NSX and vRealize capabilities on top of VxRail. It required end-to-end validation of HW and SW with interoperability and scalability testing. Further, it provides Day 2 guidance on how to monitor, backup, restore and failover management components. As such it creates a trusted implementation design that de-risks deployments, simplifies operations, and further drives IT agility for customers to create a private cloud and accelerate their transformation to a multi-cloud VMware environment. The figure below shows some of the different deployment models for modernizing vSphere and VMware environments, including the VVD using VxRail.

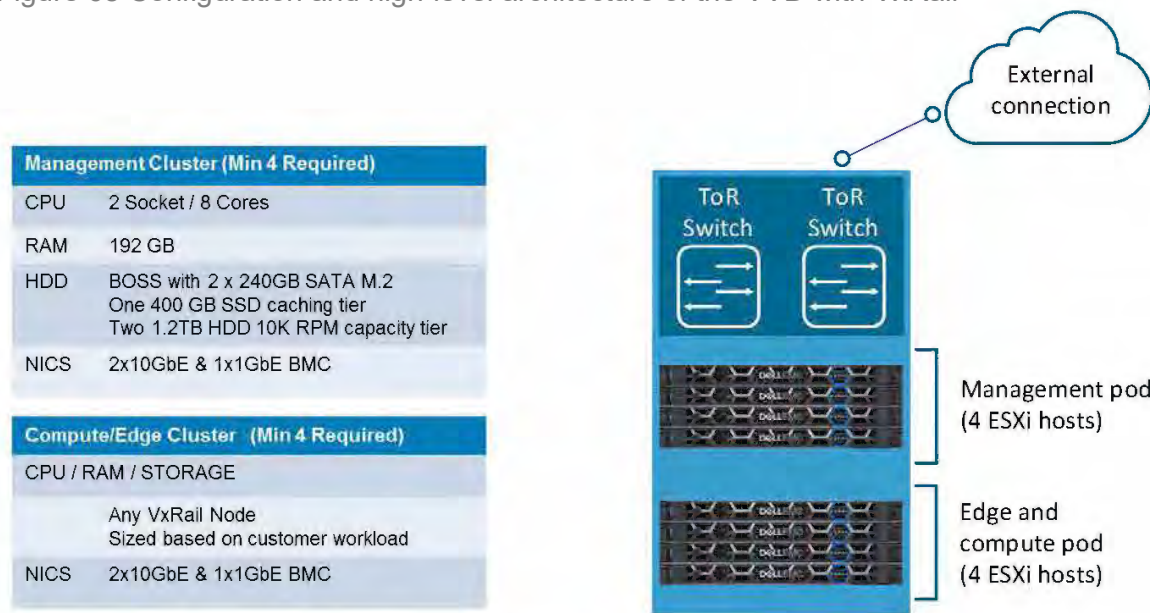
Figure 67 Paths to the VMware SDDC and multi-cloud IT model

## Accelerating IT transformation to VMware multi-cloud



The VVD requires a management cluster with four nodes and a compute/edge cluster with a minimum of 4 nodes. The compute/edge cluster can be any VxRail node as long as it adheres to the cluster design rules described in the [VxRail Scaling section](#) of this paper. The Management Cluster has a prescribed set of minimum resources. Both cluster requirements are detailed in the figure below.

Figure 68 Configuration and high-level architecture of the VVD with VxRail



Leverage the VxRail model that **meets or exceeds** these requirements and the customer needs

The VVD provides guidance to the VMware SDDC and software bill of materials as seen in the following figure. VMware recommends upgrade licenses to VMware Cloud Foundation (VCF) for VxRail customers, such that titles already paid for titles and software are not “double” charged. Purchasing VCF licenses as a bundle will result in a lower overall license cost and will also provide a path to leveraging VCF and the SDDC Manager automation more broadly on VxRail in the future. (Note: Today, SDDC Manager is not supported on VxRail. NSX and vRealize licenses included in the VCF add-on for VxRail must be run as independent titles, not as a part of SDDC Manager.)

Figure 69 Example software Bill of Materials from the VVD 4.2

Cloud Component	Product Item	Version
<b>Virtual Infrastructure</b>	ESXi	6.5 u1
	vCenter Server Appliance (VIMISO)	6.5 u1e
	NSX for vSphere	6.3.4
<b>Cloud Management</b>	vRealize Automation Appliance	7.3.0
	vRealize Orchestrator	7.3.0
	vRealize Orchestrator Plug-in for NSX	1.0.4
	vRealize Business	7.3.1
<b>Service Management</b>	vRealize Operations Manager Appliance	6.6.1
	Management Pack for NSX for vSphere	3.5.1
	Management Pack for vRealize Log Insight	6.0
	Manager Management Pack for vRealize Automation	3.0
	Management Pack for Storage Devices	6.0.5
	vRealize Log Insight	4.5.0
<b>Infrastructure</b>	Windows	2012 R2
	SQL Server	2012 R2

For more information please visit:

The VMware VVD site:

<https://www.vmware.com/solutions/software-defined-datacenter/validated-designs.html>

Dell EMC Community Network, VMware Validated Design 4.2 on VxRail Deployment Guides:

<https://community.emc.com/docs/DOC-66332>

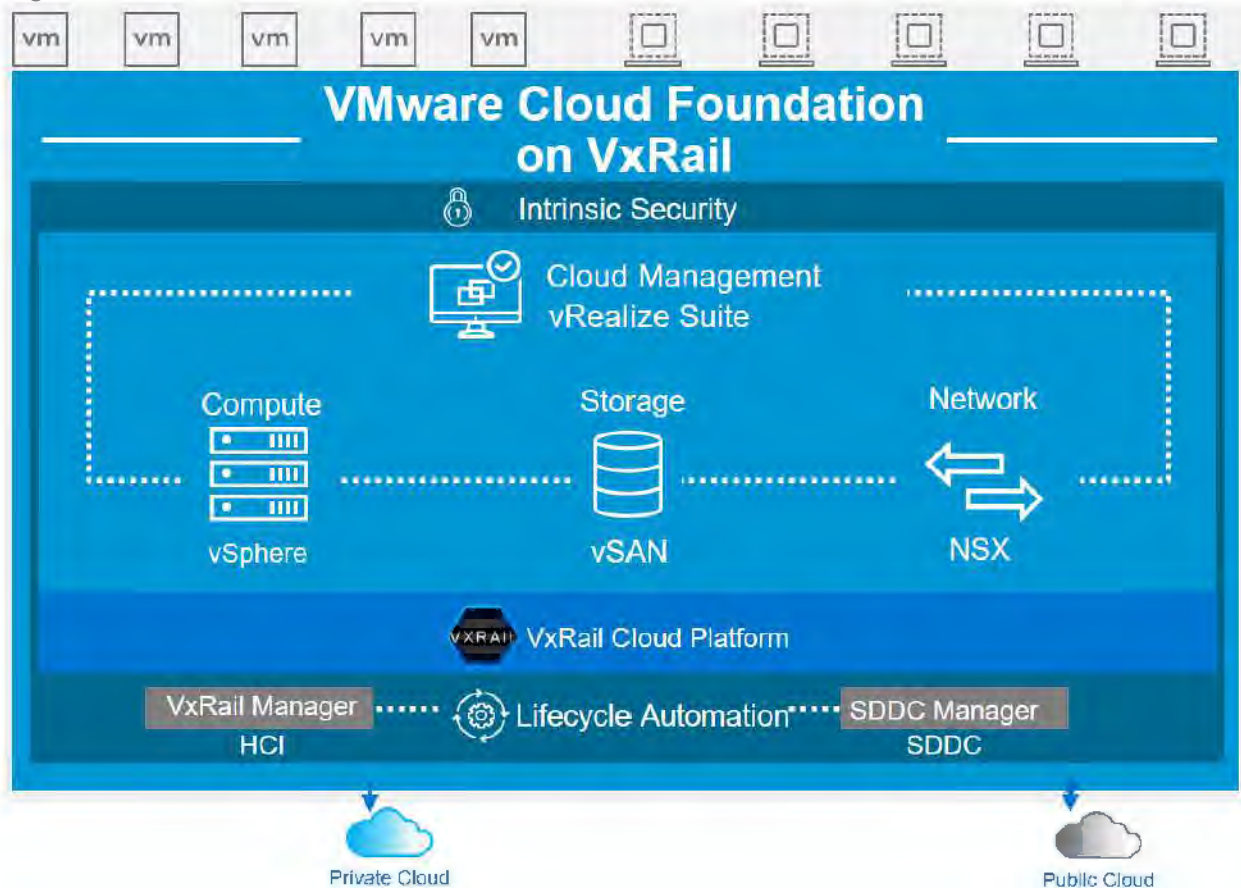
## VMware Cloud Foundation on VxRail

VMware Cloud Foundation on VxRail is a Dell EMC and VMware jointly engineered integrated solution with features that simplify, streamline, and automate the operations of your entire SDDC from Day 0 through Day 2. The new platform delivers a set of software-defined services for compute (with vSphere and vCenter), storage (with vSAN), networking and security (with NSX), and cloud management (with vRealize Suite) in both private and public environments, making it the operational hub for your hybrid cloud.

VMware Cloud Foundation on VxRail provides the simplest path to the hybrid cloud through a fully integrated hybrid cloud platform that leverages native VxRail hardware and software capabilities and other VxRail unique integrations (such as vCenter plugins and Dell EMC networking) working together to deliver a new turnkey hybrid cloud user experience with full-stack integration. Full-stack integration means you get both the HCI infrastructure layer and cloud software stack in one completely automated lifecycle turnkey experience.

An important aspect of the offering is the introduction of a standardized architecture for how these SDDC components are deployed together with the introduction of Cloud Foundation, an integrated cloud software platform that is based on VVD. Having a standardized design incorporated as part of the platform provides you with a guarantee that these components have been certified with each other and are backed by Dell Technologies. You can then be assured that there is an automated and validated path forward to get from one known good state to the next across the end-to-end stack.

Figure 70 Architecture of VMware Cloud Foundation on VxRail



## Pivotal Ready Architecture (PRA)

Pivotal Ready Architecture is a tested and validated reference architecture for deploying Pivotal Cloud Foundry on VxRail. With configurations for high availability, comprehensive product support, and options for object storage, Pivotal Ready Architecture is the best way to deploy Pivotal Cloud Foundry on-premises.

Cloud-native patterns are a modern approach to application architecture, development and delivery that has emerged as a natural response to the changes in business needs and infrastructure capabilities. This new model directly increases the speed and agility of application delivery for IT organizations and has proven its benefits for startups and established enterprises alike. Pivotal Ready Architecture is the fastest way to get Pivotal Cloud Foundry up and running in your datacenter. Accelerate your transformation with an “it just works” experience. PRA supports Pivotal Application Service (PAS) and Pivotal Container Service (PKS).

Business benefits derived from the PRA include:

- **Reliable Deployment.** PRA is a proven hardware and software solution.





- Ready Infrastructure. PRA is built on the only fully integrated, pre-configured, and pre-tested VMware hyperconverged infrastructure system family on the market.
- Resilient Architecture. PRA offers multi-site, multi-foundation, and multiple availability zone configuration options that deliver maximum uptime, geographic coverage, and resiliency.

PRA provides a tested, validated reference architecture on which to deploy a highly available enterprise-grade developer platform. Built on hyper-converged VxRail, PRA delivers automated lifecycle management of the infrastructure, a critical element in accelerating your transformation into a digital business.

- Pivotal Application Service (PAS) and Pivotal Container Service (PKS) reference architectures on VxRail
- Fully software defined infrastructure
- “Always on” highly available configurations
- A central management console
- Modular design that scales with you
- Integrated backup & disaster recovery options

For more information visit: <https://pivotal.io/pivotal-ready-architecture>

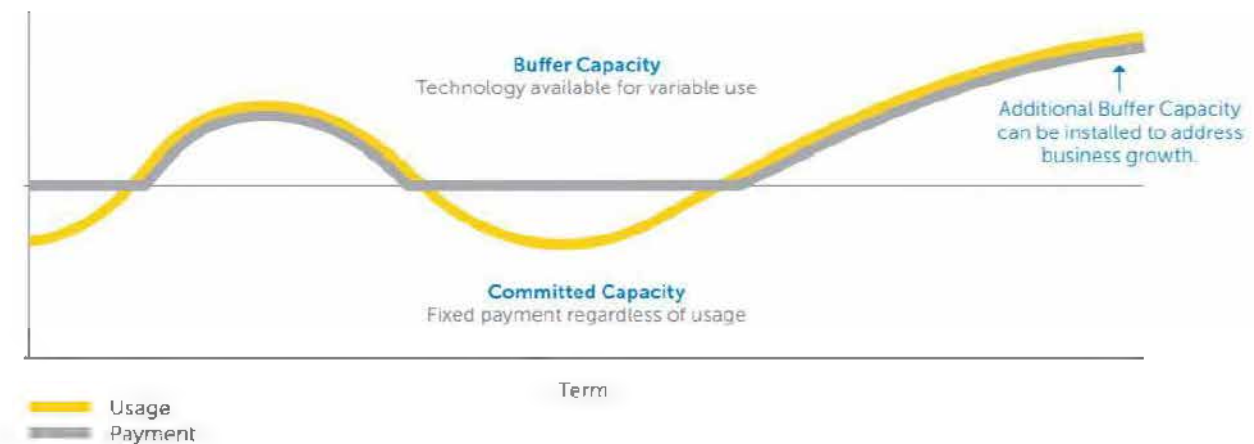
## Flexible consumption options

### Flex on Demand - a cloud-like consumption option

Flex on Demand by Dell Financial Services (DFS) allows you to acquire technology you need to support your changing business with payments that scale to match your actual usage. This model helps align your cost with usage and avoid paying for buffer capacity that is not used. It improves agility by providing instant deployment of capacity for usage when spikes occur in business operations. It improves budget agility and power by delivering better operational economics.

- DFS works with you to establish the “committed capacity” presently needed and the “buffer capacity” required in the future.
- Buffer capacity is measured using automated tools with your equipment. Each payment is comprised of the fixed committed capacity and variable buffer capacity amount.
- If your usage consistently consumes most of the buffer capacity, you have the option to receive additional buffer capacity. Once installed, your level of committed capacity and related payment will increase.

Figure 71 Relationship between technology usage and Flex on Demand payment



See <https://www.dell.com/en-us/flexibleconsumption/cloud-flex-for-hci.htm> for more information.

## VDI Complete

VDI Complete is a series of end-to-end desktop and application virtualization solutions that feature a superior solution stack and exceptional total cost of ownership. The solutions are built on Dell EMC VxRail systems and leverage VMware Horizon virtual desktops and applications.

VDI Complete is an end-to-end desktop and application virtualization solution from Dell EMC that includes everything you need to get started: the infrastructure systems, the software, the storage and the endpoints.

VDI Complete is built with best-of-breed technology from Dell Technologies. It leverages proven and trusted infrastructure systems and endpoints from Dell and Dell EMC. And it takes advantage of VMware Horizon, an industry leader desktop and application virtualization. VDI Complete is fully validated and tested, lowering risk and reducing complexity. It's a single go-to source for both purchase and complete solution support.

See [https://downloads.dell.com/solutions/general-solution-resources/White%20Papers/DellEMC.VxRail\(14G\).VMware.Horizon.RA.pdf](https://downloads.dell.com/solutions/general-solution-resources/White%20Papers/DellEMC.VxRail(14G).VMware.Horizon.RA.pdf) for more information.



**95%**

Faster initiation for 1<sup>st</sup> VxRail appliance



**47%**

Faster to provision 10 desktop VMs



**96%**

Faster to add a new appliance

## VMware Horizon

VMware Horizon is VMware's VDI and desktop-management environment. Horizon provisions user desktops using a flexible and secure delivery model. The desktop environments are accessed by the user from almost any device, including mobile devices, with the security and resiliency of the datacenter. Because the application software and data components reside in the datacenter, traditional security, backup, and disaster recovery approaches may be applied.

If a user's device is lost or the hardware fails, the recovery is straight forward. The user simply restores the environment by logging in using another device. With no data saved on the user's device, if the device is lost or stolen, there is much less chance that critical data could be retrieved and compromised.

The following figure shows how Horizon View encapsulates the OS, applications, profiles, and user data into isolated layers and dynamically assembles desktops on demand to provide users with a personalized view of their individual environments.

Figure 72 Highly available and secure desktops



Availability and security, along with ease of management and support, are compelling reasons for moving from traditional physical desktops and laptops to VDI.

VMware Horizon is a comprehensive desktop management environment that runs in a vSphere environment. The environment is managed through vCenter centralized management and can leverage advanced capabilities including, Snapshots, vMotion, DRS, and vSAN storage.

The user's desktop environment runs as a View Desktop VM on an ESXi server, and is accessed via the View Client that uses either Remote Desktop Protocol (RDP) or PC over IP protocols. The View Client can be an application running on a physical desktop, laptop, mobile device, or a web browser using the View Portal. The user's desktop environment can be either a dedicated VM or a floating VM (a VM assigned from a pool when the user logs in). Using the optional View Composer, rather than full images, linked clones can reduce the disk space required. Horizon View includes additional components used to manage the connection, provisioning the environment, authenticate users, and other applications and services.

### VMware Horizon with VxRail

The VxRail system is a self-contained compute, storage, and vSphere virtualization, and management environment that is ideally suited for VMware Horizon. VxRail accelerates the Horizon infrastructure deployment, and an environment can be up in running in hours rather than days.

VxRail hyperconverged infrastructure is available in configurations that support hundreds to thousands of virtual desktops. The number of desktops supported is based on the user-workload profile.

Dell EMC has developed tools which provide the ability to model the number of VDI environments and the expected workload profiles to determine appropriate configuration that will meet the immediate and longer term requirements. As demand increases, VxRail non-disruptively scales-up by adding additional systems and nodes while providing the users with expected performance and consistent user experience.

When deploying Horizon on VxRail systems, there are two general approaches: dedicating the VxRail environment to VDI or mixing VDI with other workloads. Horizon Editions or Horizon Add-on Editions are offered exclusively for use with VxRail. VMware or Dell EMC sales representatives can provide more details for the best customer-specific option.

In summary, VxRail with VMware Horizon allows an organization to quickly implement Desktops-as-a-Service (DaaS) and overcome the traditional capital expenditure (CAPEX) barriers of desktop virtualization. The environment can start small and easily scale up as needed. This lowers the initial startup investment. VxRail hyperconverged infrastructure is not only quick to setup, its integrated compute, storage, virtualization, and single-vendor support model eliminate the complexity of traditional infrastructure.

## VMware vSphere Platinum

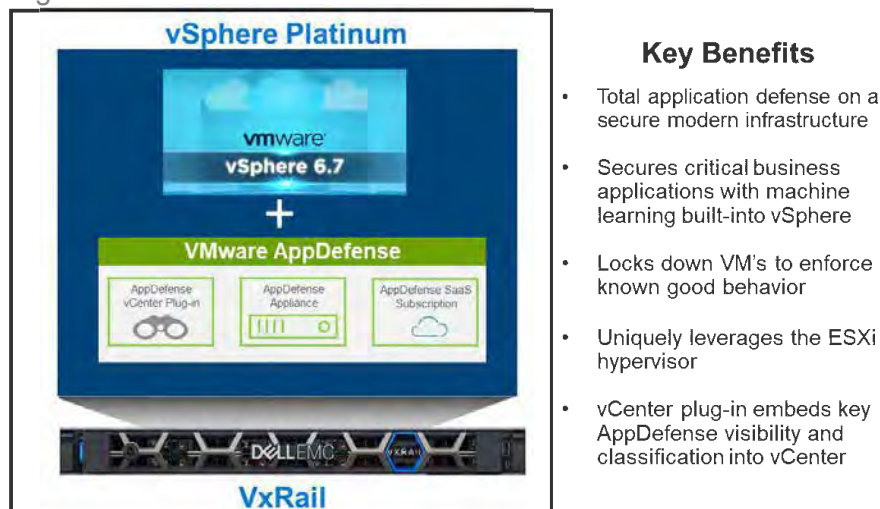
VMware vSphere® Platinum is a purpose-built security solution protecting enterprise applications, infrastructure, data, and access. It combines two proven products: vSphere, the industry-leading, efficient, and secure hybrid cloud platform for all workloads, and VMware AppDefense™, datacenter endpoint security powered by machine learning and embedding threat detection and response into the virtualization layer, to reduce security risk. While being operationally simple, vSphere Platinum ensures applications and virtual machines are running in their known-good states, with minimal overhead and performance impact. For more information about VMware vSphere Platinum, refer to

<https://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/vsphere/vmw-vsphere-platinum-solution-brief.pdf>

VMware vSphere Platinum can be run on VxRail to stand up total application defense on a secure modern infrastructure. Using Dell EMC VxRail as a platform for VMware vSphere Platinum ensures that optimized cyber resilience and security are built into every layer.

The following figure summarizes the benefits of running VMware vSphere Platinum on VxRail.

Figure 73 VMware Platinum on VxRail



### Key Benefits

- Total application defense on a secure modern infrastructure
- Secures critical business applications with machine learning built-into vSphere
- Locks down VM's to enforce known good behavior
- Uniquely leverages the ESXi hypervisor
- vCenter plug-in embeds key AppDefense visibility and classification into vCenter

## IsilonSD Edge



The EMC IsilonSD product family combines the power of Isilon scale-out NAS with the economy of software-defined storage. IsilonSD Edge is purpose-built software that addresses the need in enterprise edge locations to store growing amounts of unstructured data. IsilonSD Edge allows you to quickly deploy a simple and efficient scale-out NAS solution in a VMware environment. It also extends the reach of the data lake from your core datacenter to your edge locations by economically supporting smaller capacity deployments in a virtualized infrastructure. The data lake enables you to improve storage utilization, eliminate islands of storage and lower your TCO.

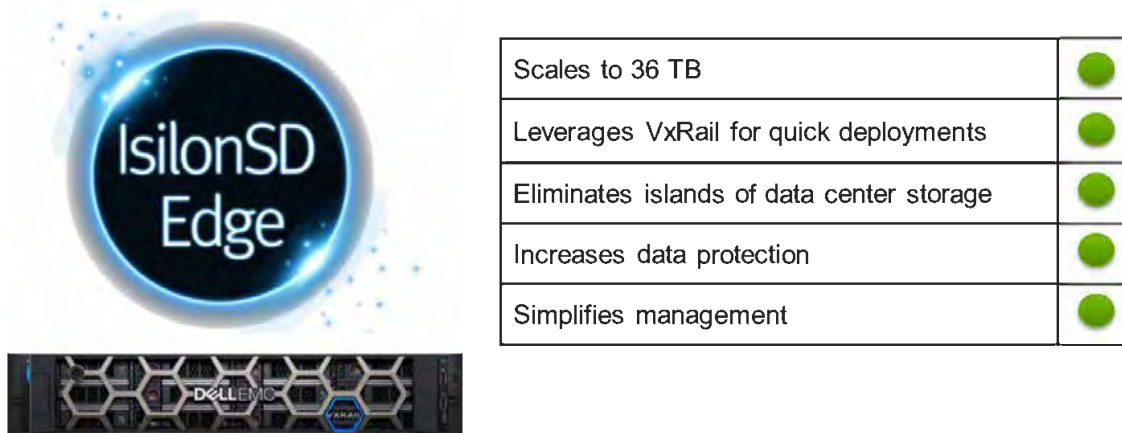
Running IsilonSD Edge on VxRail provides a simple, agile and cost-efficient platform to deliver file services from within a VxRail cluster. It is easy to manage with standard VMware tools. And it allows customers to consolidate and distribute data from and to remote locations. Best of all, it deploys in minutes. IsilonSD Edge includes all Isilon’s OneFS data services & protocols, including SMB, NFS, HDFS and OpenStack Swift.

IsilonSD Edge is tightly integrated with the VMware environment and runs on top of VMware ESXi 6.0 with VxRail 4.0 software or VMware ESXi 6.5 on VxRail 4.5 software. It leverages vCenter® with a management server that runs as a virtual image and can be used to install licenses, or add and remove nodes from a cluster. The IsilonSD Edge management server also installs a plug-in into vCenter that can be used to configure the cluster. The server and storage resources on VxRail do not need to be dedicated to IsilonSD Edge. If performance and capacity needs are met, other applications can run concurrently in the VxRail cluster. IsilonSD Edge with vSAN on VxRail is a validated and tested solution using VMFS or vSAN datastores.

The following figure summarizes the benefits of running IsilonSD Edge on VxRail.

Figure 74 IsilonSD Edge running on VxRail benefits

**Delivering Enterprise Grade File Services for VxRail**



Each IsilonSD Edge instance on your VxRail cluster can scale up to 36 TB, which is sufficient to handle the needs of many remote and branch offices. You do not have to dedicate your VxRail environment to your IsilonSD Edge cluster – you can run it alongside other workloads in the VxRail cluster.

**SAP HANA Certification with VxRail**

VxRail is among the first HCI platforms, and the first VMware-based HCI to achieve certification to run SAP HANA, SAP's in-memory database management system. SAP will leverage VxRail's persistent memory to support the application and its use cases.

Customers will benefit from running SAP HANA on VxRail because of the system's automation to get implementations up and running quickly, flexibility to offer the right mix of components to support the application from day one, and scalability to ensure future requirements are met. Start fast with automation and full lifecycle management to quickly and effectively support your HANA implementation using VxRail P Series nodes. VxRail is fully certified as a part of the Dell EMC Ready Solution for SAP v1.5 release.

VxRail is best for SAP HANA as it is fast, flexible, powerful, and scalable:

**Fast** — automation, ease of deployment / management ensure you're up and running quickly  
**Flexible** — configure a system to meet specific needs with build-to-order VxRail on PowerEdge  
**Powerful** — a rich mix of components deliver performance, density and power efficiency for both transactional process and analytics  
**Scalable** — increase power and performance without rip-and-replace system upgrades

For more information and solution guides, please visit: <https://www.dell EMC.com/en-us/solutions/business-applications/sap/hana/index.htm>

## Reference Architecture for Splunk

Splunk Enterprise is the industry-leading platform for analyzing machine-generated data. To gain valuable business insights, Splunk Enterprise uses its powerful Splunk Search Processing Language (SPL™) to extract meaningful information from machine data. The insights that are generated from analyzing machine data are called operational intelligence, which has many use cases, including:

**IT Operations**—Utilization, capacity growth

**Security**—Fraud detection, real-time detection of threats, forensics

**Internet of Things (IoT)**—Sensor data, machine-to-machine, human interactions.

Dell EMC and Splunk have partnered to provide jointly validated reference architectures that are optimized for maximum scalability and performance. Splunk software running on Dell EMC converged infrastructure delivers the operational intelligence that is required to drive an organization's digital transformation. When paired together, Dell EMC and Splunk combine the operational intelligence that is provided by the Splunk eco-system with the cost-effective, scalable, and flexible infrastructure of Dell EMC.

The primary benefits Dell EMC provides to your Splunk Enterprise environments include:

**Optimized storage data tiering**—Aligns storage to hot/warm, cold, and frozen data requirements with high retention and performance.

**Cost-effective and flexible scale-out**—Provides scale-out capacity and compute, independently or as a single, converged platform.

**Powerful data services**—Include secure encryption, compression and deduplication, and fast, efficient snapshots for protection.

A reference architecture using Dell EMC VxRail system with Isilon™ for a virtualized Splunk Enterprise environment has been jointly tested and validated by Splunk and Dell EMC to meet

or exceed the performance of Splunk Enterprise running on Splunk's documented reference hardware. VxRail offers the performance and capacity required to meet the infrastructure requirements of a small or medium-sized enterprise Splunk deployment.

See <https://www.emc.com/collateral/service-overviews/h15699-splunk-vxrail-sg.pdf> for more information.

## Additional Product information

For documentation, release notes, software updates, or for information about Dell EMC products, licensing, and service, go to the Dell EMC Online Support site (registration required) at: <https://support.emc.com>.

### Dell EMC ProSupport for Enterprise

Enterprises need unwavering support for hardware and software and a smart way to manage the mix of vendors in the datacenter. Dell EMC offers a single source with the expertise, know-how and capabilities to help you support your business.

ProSupport offers highly trained experts around the clock and around the globe to address your IT needs, minimize disruptions and maintain a high level of productivity. With over 55,000+ Dell EMC & partner professionals, across 165 countries, speaking more than 55 languages, you can rest assured that with Dell EMC you will be able to:

1. Maximize productivity by leveraging Dell EMC scale and skill
2. Minimize disruptions with around the clock access to highly trained experts
3. Gain efficiency through a single source for all your support needs

Single source, 24X7 global support is provided for VxRail system hardware and software via phone, chat, or instant message. Support also includes access to online support tools and documentation, rapid on-site parts delivery and replacement, access to new software versions, assistance with operating environment updates, and remote monitoring, diagnostics and repair with Dell EMC Secure Remote Services (ESRS).

Our 12 Centers of Excellence and Joint Solution Centers deliver in-house collaboration and industry-leading levels of support, leveraging Dell EMC's alliances with leading application providers such as Oracle and Microsoft. Our 87 technical support sites are comprised of 71 total Dell Tech Support Sites and 16 total EMC Customer Service Centers.

Dell EMC support is recognized with 94% customer satisfaction rating and has received multiple awards including Temkin Group CE Excellence, TSIA STAR awards, Microsoft Deployment Partner of the Year and many more.

The Dell EMC difference is clear, when it comes to your IT strategy we allow you to fearlessly adopt new technology giving you freedom to focus on your business. Having the same enterprise-class support from Dell EMC across your infrastructure gives you that freedom.

### Dell EMC ProDeploy Services for VxRail Systems

Dell EMC offers ProDeploy installation and implementation services to ensure smooth and rapid integration of VxRail systems into customer networks. The standard service, optimal for a single system, provides an expert on site to perform a pre-installation checklist with the data-center team, confirm the network and Top of Rack (TOR) switch settings, conduct site validation, rack and cable, configure, and initialize the system. Finally, an on-site Dell EMC service technician will configure EMC Secure Remote Services (ESRS) and conduct a brief functional overview on essential VxRail system administrative tasks. A custom version of this installation and implementation service is available for larger-scale VxRail system deployments, including those with multiple systems or clustered environments. Also offered is VxRail system extended service, which is delivered remotely and provides an expert service technician to rapidly

implement VxRail system pre-loaded data services (RecoverPoint for Virtual Machines, and vSphere Data Protection).



# APPLIANCE DELL EMC VXRAIL™

na 14ª geração dos servidores Dell EMC PowerEdge

## O padrão em infraestrutura hiperconvergente

O appliance Dell EMC VxRail™, o exclusivo dispositivo de infraestrutura hiperconvergente da Dell EMC e da VMware, é a maneira mais rápida e fácil de ampliar e simplificar um ambiente VMware. Habilitado pelo VMware vSAN™ e gerenciado por meio da interface do vCenter, o appliance Dell EMC VxRail oferece aos clientes existentes da VMware uma experiência com a qual já estão familiarizados. A integração perfeita com as ferramentas existentes da VMware também permite que os clientes aproveitem e expandam suas atuais ferramentas e processos de TI.

A arquitetura do appliance Dell EMC VxRail é um sistema distribuído que consiste em componentes modulares comuns, que podem ser dimensionados linearmente de 3 a 64 nós em um cluster. Com o poder de toda a SAN (Storage Area Network), ele oferece uma solução hiperconvergente simples e econômica que fornece múltiplas opções de computação, memória, armazenamento, rede e gráficos para corresponder a qualquer caso de uso e cobrir uma ampla variedade de aplicativos e cargas de trabalho.

Com base no VMware vSAN e no vSphere, os programas de software líderes do setor, e desenvolvido com os novos processadores dimensionáveis Intel® Xeon®, o appliance Dell EMC VxRail permite que os clientes comecem pequeno e cresçam, dimensionando capacidade e desempenho facilmente e de modo não disruptivo. O dimensionamento de único nó e a expansão da capacidade de armazenamento proporcionam uma abordagem previsível, do tipo "pague conforme crescer", para crescimento futuro conforme necessário. Desenvolvido com base na 14ª geração dos servidores PowerEdge, a base do datacenter, o VxRail é projetado para as cargas de trabalho de missão crítica dos dias atuais, oferecendo mais processador, armazenamento flash e opções de conectividade de rede de que seu antecessor. O desempenho está melhor do que nunca, com tempos de resposta 2x melhores e até 2x mais IOPS (Input/Output per Second).

Serviços de dados de missão crítica acompanham o appliance Dell EMC VxRail sem nenhum custo adicional. Tecnologias de proteção de dados, inclusive EMC RecoverPoint for VMs e VMware vSphere Data Protection, estão incorporadas ao appliance, com a opção de adicionar o Data Protection Suite for VMware e o DD VE (Data Domain Virtual Edition) para ambientes maiores que exigem proteção de dados mais abrangente.

O appliance Dell EMC VxRail também tem o respaldo do suporte de nível internacional da Dell EMC, com um único ponto de contato para hardware e software, e inclui o Dell EMC ESRS para dial home e conexão remota bidirecional proativa para monitoramento remoto, diagnóstico e reparo a fim de garantir disponibilidade máxima.

As especificações detalhadas e um comparativo dos appliances Dell EMC VxRail na 14ª geração dos servidores PowerEdge é exibida abaixo.

	Série E	Série V	Série P	Série S
<b>Computação, armazenamento e memória (por nó)</b>				
Chassis	1U1N	2U1N	2U1N	2U1N
Família de processadores dimensionáveis Intel™ Xeon™				
Soquetes da CPU	Simple ou duplos	Duplos	Simple ou duplos	Simple ou duplos
Núcleos de CPU	4 a 56	8 a 56	8 a 56	4 a 56
Frequência da CPU	2,0 GHz a 3,6 GHz	2,0 GHz a 3,6 GHz	2,0 GHz a 3,6 GHz	2,0 GHz a 3,6 GHz
RAM*	96 GB a 1.536 GB	192 GB a 1.536 GB	96 GB a 1.536 GB	96 GB a 1.536 GB
SSD do cache	400 GB a 1.600 GB	400 GB a 1.600 GB	400 GB a 1.600 GB	400 GB a 1.600 GB
Armazenamento híbrido	1,2 TB a 16 TB	1,2 TB a 40 TB	1,2 TB a 40 TB	4 TB a 48 TB
Armazenamento totalmente flash (SAS ou SATA)	1,92 TB a 30,7 TB	1,92 TB a 76,8 TB	1,92 TB a 76,8 TB	Somente híbrido
Gabinetes de unidades	10 x 2,5 pol.	24 x 2,5 pol.	24 x 2,5 pol.	12 x 3,5 pol. e 2 x 2,5 pol.
Máx. de grupos de disco	2	4	4	2
Solução de armazenamento otimizada para inicialização (BOSS)	2 x 240 GB SATA M.2 RAID 1	2 x 240 GB SATA M.2 RAID 1	2 x 240 GB SATA M.2 RAID 1	2 x 240 GB SATA M.2 RAID 1
Máx. de GPUs PCIe	n/d	1x a 2x NVIDIA Tesla M10 ou 1x a 3x NVIDIA Tesla M60	n/d	n/d

\*Valores além de 768 GB de RAM exigem CPUs de soquete duplo

\*\*SSD de cache de 1.600 GB apenas em configurações híbridas

\*\*\*A adição de GPUs reduz o total de RAM e a conectividade de rede adicional

	Série E	Série V	Série P	Série S
<b>Clustering e dimensionamento</b>				
Máx. de nós* (por cluster)	64	64	64	64
Mín. de nós (por cluster)	3	3	3	3
Incrementos de dimensionamento (em nós)	1	1	1	1

\*Máximo de 8 nós por cluster em modelos de 1 GbE


	Série E	Série V	Série P	Série S
<b>Sistema de rede (por nó)</b>				
Conectividade do appliance	4 x RJ45 de 10 GbE ou 4 x SFP+ de 10 GbE ou 4 x RJ45 de 1 GbE*	4 x RJ45 de 10 GbE ou 4 x SFP+ de 10 GbE	4 x RJ45 de 10 GbE ou 4 x SFP+ de 10 GbE ou 4 x RJ45 de 1 GbE*	4 x RJ45 de 10 GbE ou 4 x SFP+ de 10 GbE ou 4 x RJ45 de 1 GbE*
Portas de gerenciamento	1 x RJ45 iDRAC9 Enterprise de 1 GbE	1 x RJ45 iDRAC9 Enterprise de 1 GbE	1 x iDRAC9 de 1 GbE Enterprise RJ45	1 x iDRAC9 de 1 GbE Enterprise RJ45
Conectividade opcional (máx. de portas adicionais)	Até 8 x RJ45 de 10 GbE	Até 16 x RJ45 de 10 GbE ou Até 16 x SFP+ de 10 GbE	Até 16 x RJ45 de 10 GbE ou Até 16 x SFP+ de 10 GbE	Até 12 x RJ45 de 10 GbE ou Até 12 x SFP+ de 10 GbE

\*Conectividade de 1 GbE limitada à CPU de soquete simples



	Série E	Série V	Série P	Série S
<b>Alimentação e dimensões</b>				
PSU redundante dupla de alta eficiência	1.100 W 100V a 240V AC 1.100 W - 48V CC	2.000 W 200V a 240V AC	1.100 W 100V a 240V AC 1.100 W - 48V CC 1.600 W 200V a 240 V AC	1.100 W 100V a 240V AC 1.100 W - 48V CC
Ventiladores de refrigeração redundantes	8	6	4 ou 6	6
Dimensões físicas	42,8 mm/1,68 pol. A 434,0 mm/17,09 pol. L 733,82 mm/29,61 pol. P 21,9 kg/48,28 lb.	86,8 mm/3,42 pol. A 434 mm/17,09 pol. L 678,8 mm/26,72 pol. P 28,1 kg/61,95 lb.	86,8 mm/3,42 pol. A 434 mm/17,09 pol. L 678,8 mm/26,72 pol. P 28,1 kg/61,95 lb.	86,8 mm/3,42 pol. A 434 mm/17,09 pol. L 678,8 mm/26,72 pol. P 33,1 kg/72,91 lb.

	Série E	Série V	Série P	Série S
<b>Ambiental e certificações</b>				
Temperatura operacional ambiente	10 °C a 30 °C 50 °F a 86 °F	10 °C a 30 °C 50 °F a 86 °F	10 °C a 30 °C 50 °F a 86 °F	10 °C a 25 °C 50°F a 77°F
Faixa de temperatura de armazenamento	-40 °C a +65 °C -40 °F a +149 °F	-40 °C a +65 °C -40 °F a +149 °F	-40 °C a +65 °C -40 °F a +149 °F	-40 °C a +65 °C -40 °F a +149 °F
Umidade relativa operacional	10% a 80% (sem condensação)	10% a 80% (sem condensação)	10% a 80% (sem condensação)	10% a 80% (sem condensação)
Altitude de operação sem reduções de taxa	3.048 m aprox. 10.000 pés	3.048 m aprox. 10.000 pés	3.048 m aprox. 10.000 pés	3.048 m aprox. 10.000 pés
Dissipação de calor	4.100 BTU/h	7.500 BTU/h	6.000 BTU/h	4.416 BTU/h
Certificações	UL (Cobre cUL e não exige CSA), CE, EMC, FCC			



Saiba mais sobre os appliances Dell EMC VxRail



Entre em contato com um especialista da Dell EMC

0800 553622

# vSAN™ 2-Node Cluster on VxRail™ Planning Guide

## Abstract

This guide provides information for the planning of a VMware vSAN 2-Node Cluster infrastructure on a VxRail platform. This guide will focus on the VxRail implementation of the vSAN 2-Node Cluster, including minimum requirements and recommendations.

May 2019

# Table of contents

1.0 Overview .....	3
1.1 INTRODUCTION .....	4
2.0 Requirements, Recommendations, and Restrictions .....	5
2.1 VXRAIL HARDWARE .....	5
2.2 VXRAIL SOFTWARE VERSION .....	5
2.3 VMWARE VCENTER SERVER .....	5
2.4 WITNESS VIRTUAL APPLIANCE .....	5
Software version .....	5
Installation .....	5
Sizing .....	6
2.5 PHYSICAL NETWORK .....	6
2.6 PORT REQUIREMENTS .....	7
2.7 WITNESS AND MANAGEMENT NETWORK TOPOLOGY .....	7
2.8 NETWORK LAYOUT .....	8
2.9 CAPACITY PLANNING CONSIDERATIONS .....	9
Storage Capacity .....	9
CPU & Memory Capacity .....	9
Network Bandwidth .....	10
2.10 UPGRADE OPTIONS .....	11
2.11 LICENSING .....	12
3.0 Deployment Types .....	13
OPTION 1: CENTRALIZED MANAGEMENT .....	13
OPTION 2: CENTRALIZED MANAGEMENT, LOCALIZED WITNESS .....	13
OPTION 3: LOCALIZED MANAGEMENT AND WITNESS .....	14
4.0 Conclusion .....	15
5.0 References .....	16

## 1.0 Overview

VMware vSAN 2-Node Cluster is a configuration implemented in environments where a minimal configuration is a key requirement, typically in Remote Office and Branch Office (ROBO) such as retail stores.

The VxRail 4.7.100 is the first release to support the vSAN 2-Node Cluster Direct Connect configuration.

This guide provides information for the planning of a vSAN 2-Node Cluster infrastructure on a VxRail platform. This guide will focus on the VxRail implementation of the vSAN 2-Node Cluster, including minimum requirements and recommendations.

For detailed information about VMware vSAN 2-Node Clusters architecture and concepts, please refer to the [VMware vSAN 2-Node Guide](#).

## 1.1 INTRODUCTION

A VMware vSAN 2-Node Cluster on VxRail consists of a cluster with two directly connected VxRail E560 or E560F nodes, and a Witness Host deployed as a Virtual Appliance. The VxRail cluster is deployed and managed by VxRail Manager and VMware vCenter Server™.

A vSAN 2-Node configuration is very similar to a Stretched Cluster configuration. The Witness Host is the component that provides quorum for the two data nodes in the event of a failure. As in a stretched cluster configuration, the requirement for one Witness per cluster still applies.

Unlike a Stretched Cluster, typically the vCenter Server and the Witness Host are located in a main datacenter, as illustrated below, and the two vSAN data nodes are in a remote location. Even though the Witness host can be deployed at the same site as the data nodes, the most common deployment for multiple 2-node clusters is to have multiple Witnesses hosted in the same management cluster as the vCenter Server, optimizing the infrastructure cost by sharing the vSphere licenses and the management hosts.

This design is facilitated by the low bandwidth required for the communication between data nodes and the Witness.

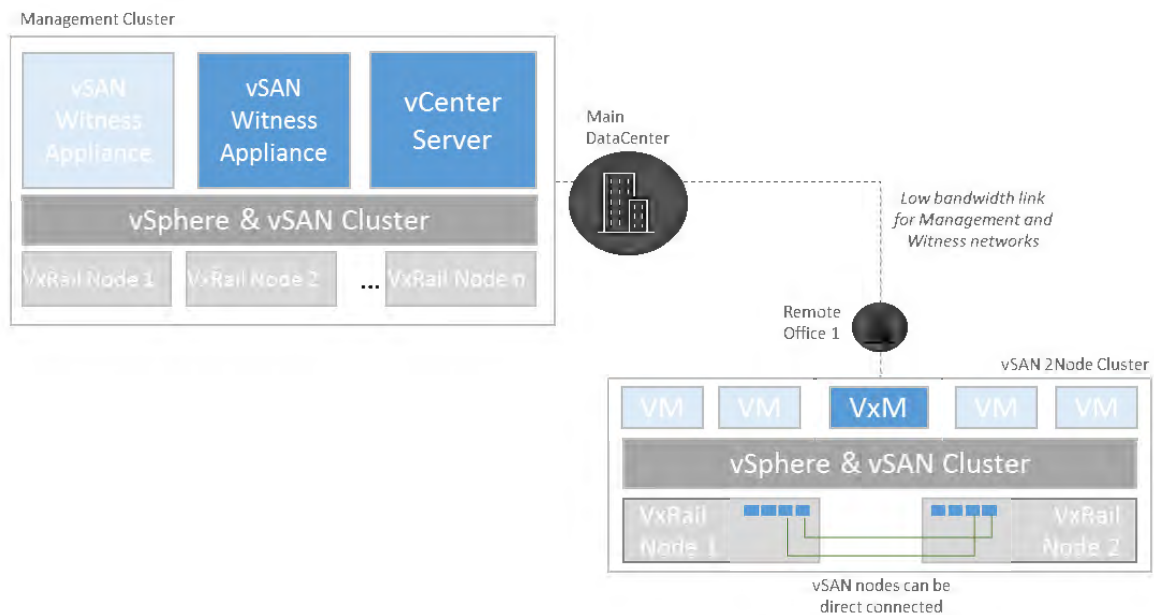


Figure 1: Design of 2-Node Cluster with Witness hosted in a centralized datacenter

A vSAN 2-Node configuration maintains the high availability characteristics as a regular cluster. Each physical node is configured as a vSAN Fault Domain. This means the virtual machines can have one copy of data on each fault domain. In the event of a node or a device failure, the virtual machine remains accessible through the alternate replica and Witness components.

When the failed node is restored, the Distributed Resource Scheduler (DRS) automatically rebalances the virtual machines between the two nodes. DRS is highly recommended and it requires a vSphere Enterprise edition license or higher.

## 2.0 Requirements, Recommendations, and Restrictions

### 2.1 VXRAIL HARDWARE

In VxRail 4.7.100, the VxRail E-Series models E560 and E560F. The systems can be configured with the following Network Daughter Card.

- 4 x 10GbE



Figure 2: Front and back views of the VxRail Appliance

### 2.2 VXRAIL SOFTWARE VERSION

VxRail 4.7.100 or later is required.

### 2.3 VMWARE VCENTER SERVER

The vSAN 2-Node Cluster must be connected to an external vCenter Server at the time of its deployment.

- VMware vCenter Server version 6.7u1 is the minimum required.
- The vCenter Server must be deployed before the deployment of the 2-Node Cluster.
- vCenter Server cannot be deployed on the 2-Node Cluster.

### 2.4 WITNESS VIRTUAL APPLIANCE

VMware supports both physical ESXi Hosts and Virtual Appliance as vSAN Witness Host. VxRail 4.7.100 will only support using the vSAN Witness Virtual Appliance. The Witness Virtual Appliance does not consume extra vSphere licenses and does not require a dedicated physical host.

#### Software version

- vSAN Witness Appliance version 6.7u1 is the minimum requirement.
- Witness Appliance must be at the same version as the ESXi hosts.
- The vSphere license is included and hard-coded in the Witness Virtual Appliance.

#### Installation

- The Witness Appliance must be installed, configured, and added to vCenter inventory before the vSAN 2-Node Cluster on VxRail deployment.
- The Witness Appliance must have connectivity to both vSAN nodes.
- The Witness Appliance must be managed by the same vCenter Server that is managing the 2-Node Cluster.
- A Witness Appliance can only be connected to one vSAN 2-Node Cluster.

- The general recommendation is to place the vSAN Witness Host in a different datacenter, such as a main datacenter or a cloud provider.
- The Witness can run in the same physical site as the vSAN data nodes but, cannot be placed in the 2-Node cluster to which it provides quorum.
- It is possible to deploy the Witness Appliance on another 2-Node Cluster, but it is not recommended. A VMware RPQ is required for this solution design.

#### Sizing

- There are 3 typical sizes for a witness appliance that can be selected during deployment: Tiny, Normal and Large. Each option has different requirements for compute, memory and storage.

##### *Tiny (10 VMs or fewer)*

- 2 vCPUs, 8 GB vRAM
- 8 GB ESXi Boot Disk, one 10 GB SSD, one 15 GB HDD
- Supports a maximum of 750 witness components

##### *Normal (up to 500 VMs)*

- 2 vCPUs, 16 GB vRAM
- 8 GB ESXi Boot Disk, one 10 GB SSD, one 350 GB HDD
- Supports a maximum of 22,000 witness components

##### *Large (more than 500 VMs)*

- 2 vCPUs, 32 GB vRAM
- 8 GB ESXi Boot Disk, one 10 GB SSD, three 350 GB HDDs
- Supports a maximum of 45,000 witness components

#### *Figure 3: Sizing guidance from VMware vSAN planning guide*

- The general recommendation is to use the normal size. However, 2-Node clusters with up to 25 VMs are good candidates for the “Tiny” option because they are less likely to reach or exceed 750 components.
  - Each storage object is deployed on vSAN as a RAID tree and each leaf of the tree is said to be a component. For instance, when we deploy a VMDK with a RAID-1 mirror, we will have a replica component in one host and another replica component on another host. The number of stripes used has an effect, i.e., if using 2 stripes we will have 2 replica components in each host.

## 2.5 PHYSICAL NETWORK

In the VxRail 4.7.100 release, the two vSAN data nodes must be direct connected using a network crossover cable, or SFP+ cables.

A physical layout is enforced.

- Either a 1GbE or a 10GbE switch is supported.
- Ports 1 and 2 of the VxRail appliances are connected to a switch and used for the management and witness traffic. Port speed will auto-negotiate down to 1Gb if connected to a 1GbE switch.




Figure 4: Port configuration on VxRail appliances

- Ports 3 and 4 from Node 1 are direct connected to Ports 3 and 4 of Node 2 respectively and are used for vSAN and vMotion traffic.

Because the two VxRail nodes are direct connected, the latency between the nodes is within the recommended 5msec roundtrip time (<2.5ms one-way).

## 2.6 PORT REQUIREMENTS

The list below is for services that are needed. The incoming and outgoing firewall ports for these services should be opened.

Services	Port #	Protocol	To/From
vSAN Clustering Service	12345,2345	UDP	vSAN Hosts
vSAN Transport	2233	TCP	vSAN Hosts
vSAN VASA Vendor Provider	8080	TCP	vSAN Hosts & vCenter Server
vSAN Unicast Agent to the Witness Host	12321	UDP	vSAN Hosts & Witness Appliance

Figure 5: Service ports on VxRail appliance

## 2.7 WITNESS AND MANAGEMENT NETWORK TOPOLOGY

VMware recommends that the vSAN communications between vSAN Nodes and the vSAN Witness Host be:

- Layer 2 (same subnet) for configurations with the Witness Host in the same location
- Layer 3 (routed) for configurations with the Witness Host in an alternate location such as at the main datacenter
  - A static route is required

The maximum supported roundtrip time (RTT) between the vSAN 2-Node Cluster and the Witness is 500ms (250ms each way).

In the VxRail implementation of the vSAN 2-Node Cluster, a VMkernel interface is designated to carry traffic destined for the Witness Host.



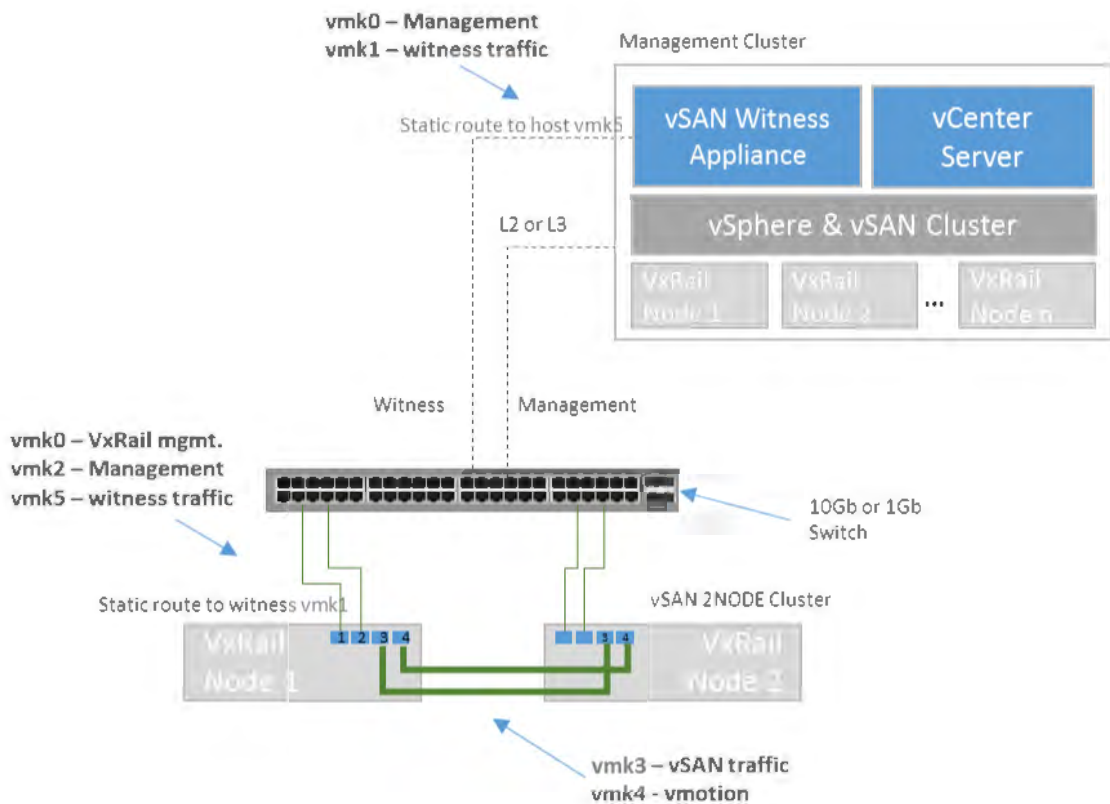


Figure 6: Port configuration for traffic between 2-Node Cluster and Witness Host

Each vSAN Host's vmk5 VMkernel interface is tagged with "witness" traffic. When using layer 3, each vSAN Host must have a static route configured for vmk5, able to properly access the vmk1 on the vSAN Witness Host, which is tagged with "vsan" traffic.

Likewise, the vmk1 interface on the witness host must have a static route configured to properly communicate with vmk5 on each vSAN Host.

## 2.8 NETWORK LAYOUT

The chart below illustrates the network layout used by VxRail in the configuration of a vSAN 2-Node Cluster. One additional VLAN is needed for Witness Traffic Separation. This layout is specific to the VxRail vSAN 2-Node Cluster. The configuration of the management cluster will be slightly different as described in the VxRail Networking Guide.

Traffic	Requirements	Members	NIOC Shares	UPLINK1	UPLINK2	UPLINK3	UPLINK4
				VMNIC0	VMNIC1	VMNIC2	VMNIC3
Management Network	Same VLAN as vCenter Server network	ESXi vmk2,	40	Standby	Active	Unused	Unused
vCenter Server Network	Same VLAN as management network	vCenter, VxM VNIC0				Unused	Unused
VxRail Management	VLAN 3939, can't change the physical port.	ESXi vmk0, VxM VNIC1(New)		Standby	Active	Unused	Unused
Witness Traffic		ESXi vmk5		Active	Standby		
vSAN		ESXi vmk3	100	Unused	Unused	Active	Standby
vMotion		ESXi vmk4	50	Unused	Unused	Standby	Active
Virtual Machines		Virtual Machines	60	Active	Standby	Unused	Unused

Figure 7: Network layout of a VxRail 2-Node Cluster

## 2.9 CAPACITY PLANNING CONSIDERATIONS

In this section we offer general recommendations for storage, CPU, memory, and link bandwidth sizing.

### Storage Capacity

- The minimum of 25% to 30% of spare storage capacity remains an adequate requirement for a 2-Node Cluster.
- Note that in a 2-Node Cluster, the protection method will be RAID-1 and in case of a node failure the surviving node will continue to operate with a single object's component.

### CPU & Memory Capacity

- When defining CPU and Memory capacity, consider the minimum capacity needed to satisfy the VM requirements while in a failed state.
- The general recommendation is to size a cluster to operate below 50% of the max CPU required, taking in consideration the projected growth in consumption.

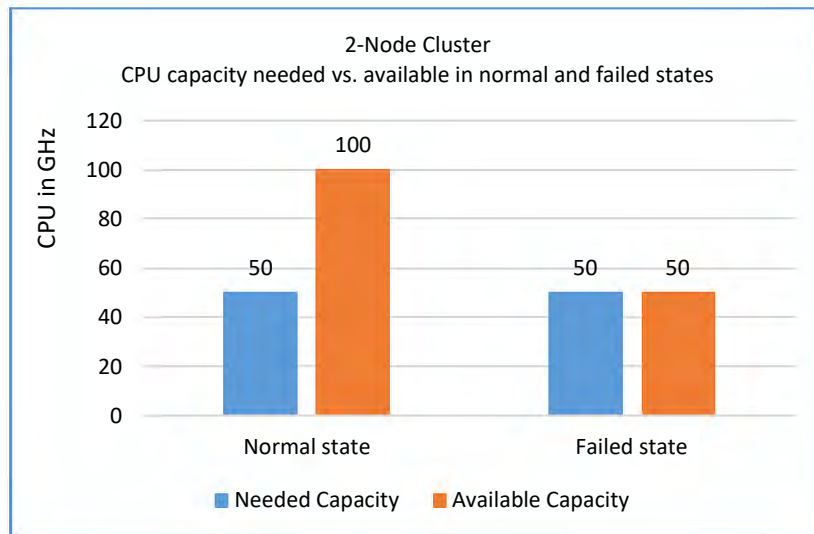


Figure 8: CPU capacity planning

### Network Bandwidth

Our measurements indicate a regular T1 link can satisfy the network bandwidth requirements for the communications between Data Nodes <> vCenter Server and Data Nodes <> Witness Appliances.

However, with the purpose of adapting the solution to different service level requirements, it is important to understand in more details the requirements for:

- Normal cluster operations
- Witness contingencies
- Services, such as maintenance, lifecycle management, and troubleshooting

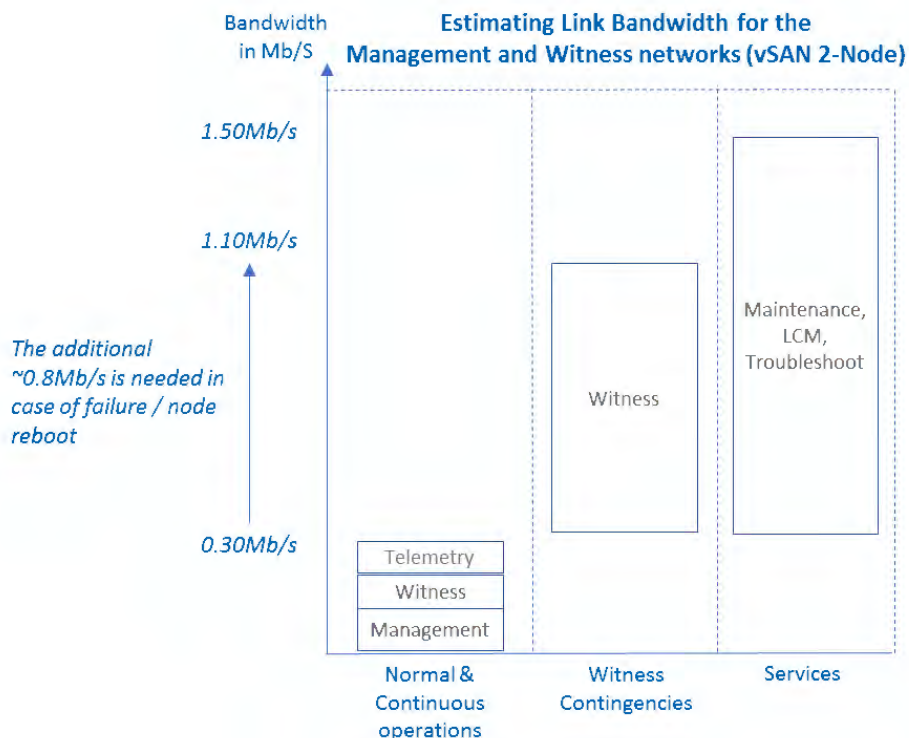


Figure 9: Network bandwidth planning considerations

#### *Normal Cluster Operations*

- Normal cluster operations include the traffic between data nodes, vCenter Server and the Witness Appliance.
- During normal operations, the bulk of the traffic is between data nodes and vCenter Server. This traffic is affected primarily by number of VMs and number of components but, is typically very light load.
- Our measurements of a cluster with 25 VMs and near 1000 components indicated a bandwidth consumption lower than 0.3Mbps

#### *Witness Contingencies*

- The Witness Appliance does not maintain any data, only metadata component.
- The Witness traffic can be influenced by the IO workload running in the cluster, but in general this is very small traffic while the cluster is in a normal state.
- However, in the event the preferred Witness Host fails or is partitioned,
  - vSAN powers off the VMs in the failed host
  - The secondary node is elected as the HA master and the Witness Host sends updates to the new master, which are followed by the acknowledgement from the master that the ownership is updated
  - 1138 bytes are required for each component update
  - When the update is completed, quorum is formed between the secondary host and the Witness Host, allowing the VMs to have access to their data and be powered on.
- The failover procedure requires enough bandwidth to allow for the ownership of components to change within a short interval of time.
- Our recommendation for a 2-Node Cluster with up to 25 VMs is to ensure that at least 0.8Mbps is available to ensure a successful failover operation.

#### *Maintenance, Lifecycle Management & Troubleshooting*

- The amount of bandwidth reserved for maintenance, lifecycle management and troubleshooting are determined primarily by the desired transfer times for large files.
- The log files used in troubleshooting are compressed and typically can be transferred in a reasonable time.
- However, the composite files used for software and firmware upgrades can be up to 4.0GB and can take a long time to be transferred when using a T1 link. The bandwidth requirements should be evaluated in case the customer has specific maintenance window requirements.
  - As a reference, if using a T1 link, we will expect that at least 1Mb/s of bandwidth will be available for the transfer of the composite file and we can estimate that this transfer will happen in about 9 hours.

## 2.10 UPGRADE OPTIONS

VxRail supports two options for node upgrades:

- 1) Fully automated
  - a. All components including Witness nodes are upgraded via VxRail LCM
- 2) Witness manual upgrade
  - a. Customers manually upgrade witness nodes

- b. All other components are auto upgraded via VxRail LCM

## 2.11 LICENSING

Any of the licensing editions can be used on a vSAN 2-Node Cluster.

	Standard	Advanced	Enterprise
Storage Policy Based Mgmt.	✓	✓	✓
Virtual Distributed Switch	✓	✓	✓
Rack Awareness	✓	✓	✓
Software Checksum	✓	✓	✓
All-Flash Hardware	✓	✓	✓
iSCSI Target Service	✓	✓	✓
QoS - IOPS Limit	✓	✓	✓
Deduplication & Compression		✓	✓
RAID-5/6 Erasure Coding		✓	✓
vRealize Operations within vCenter		✓	✓
Data-at-Rest Encryption			✓
Stretched Cluster with Local Failure Protection			✓

Figure 10: Detailed chart of vSphere licensing options

For more information, see the [VMware vSAN Licensing guide](#).

Please note that VxRail 4.7.100 does not support expansion to more than two nodes so some of the features included in the license edition are not be available.

Witness Appliance license is not required but the host where the Witness resides needs the appropriate vSphere license.

### 3.0 Deployment Types

#### OPTION 1: CENTRALIZED MANAGEMENT

In this scenario, customer vCenter servers, and Witness virtual appliances are deployed at the same management cluster located at a main datacenter. One vCenter server instance can manage multiple VxRail vSAN 2-Node Clusters but, each VxRail vSAN 2-Node Cluster must have its own Witness.

Network bandwidth must be within the minimum requirement as stated earlier. Enhanced Link Mode is recommended.

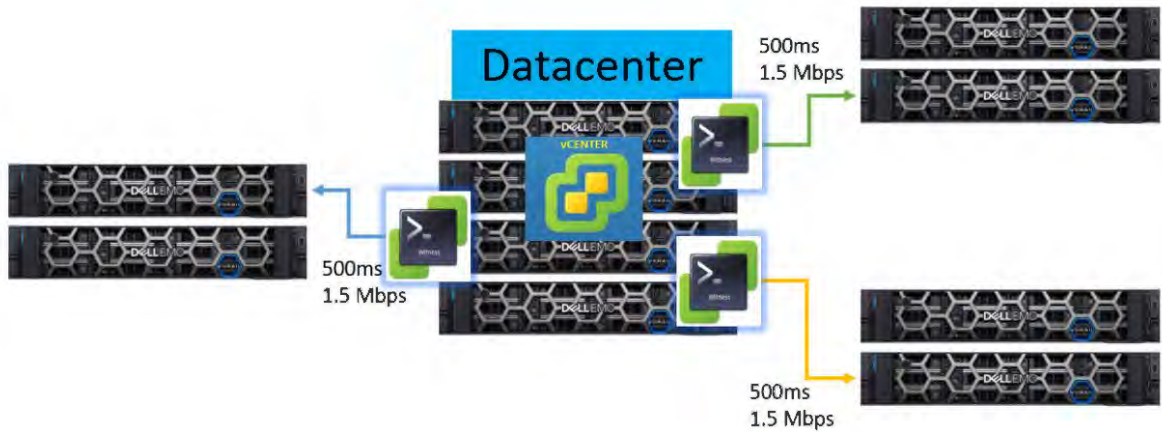


Figure 11: Centralized management of vCenter Server and Witness Appliances

#### OPTION 2: CENTRALIZED MANAGEMENT, LOCALIZED WITNESS

In this deployment option the vCenter server is located at the main datacenter, but the vSAN Witness Appliance and the two VxRail nodes are at the same location. An additional ESXi host is required to host vSAN Witness Appliance. vSAN Witness Appliance cannot be hosted in the VxRail 2-Node Cluster.




Figure 12: Centralized vCenter with local Witness Appliances



### OPTION 3: LOCALIZED MANAGEMENT AND WITNESS

In this option, the three fault domains are at the same location; the vCenter server, vSAN Witness Appliance and the VxRail Nodes. An additional ESXi host is required to host vSAN Witness Appliance and customer vCenter Server. vSAN Witness Appliance and customer-supplied vCenter cannot be hosted in the VxRail vSAN 2-Node Cluster.



Figure 13: Localized vCenter and Witness Appliance

### Considerations about the deployment options

Option	Pros	Cons
Centralized Management and Witness	<ul style="list-style-type: none"> <li>- Single pane of glass for the management of multiple 2-Node Clusters</li> <li>- Centralization of Witness appliances reduces licensing and hardware costs</li> </ul>	<ul style="list-style-type: none"> <li>- Network costs for vCenter and Witness communications</li> </ul>
Centralized Management, localized Witness	<ul style="list-style-type: none"> <li>- Single pane of glass for the management of multiple 2-Node Clusters</li> </ul>	<ul style="list-style-type: none"> <li>- Network costs for vCenter communications</li> <li>- Software and hardware costs for deployment of witness appliances</li> </ul>
Localized management and witness	<ul style="list-style-type: none"> <li>- Reduces network cost associated to normal operations and witness contingency</li> </ul>	<ul style="list-style-type: none"> <li>- Software and hardware costs for deployment of multiple vCenter Servers and witness appliances</li> <li>- Network Bandwidth still needed for maintenance and troubleshooting which is the larger bandwidth requirement</li> </ul>

#### 4.0 Conclusion

Starting with VxRail 4.7.100, VMware vSAN 2-Node Cluster direct connect is supported using E560/F Dell PowerEdge platform. A VMware vSAN 2-Node Cluster is a minimal configuration consisting of two vSAN data nodes and a Witness Virtual Appliance.

vSAN 2-Node Cluster can easily be deployed anywhere, but mainly targeted at Remote Offices and Branch Offices (ROBO). Many vSAN 2-Node Clusters can be managed by a single vCenter instance. This minimal configuration continues to provide the same functional benefits of vSphere and vSAN. It enables an efficient centralized management with reduced hardware and software costs, fitting well the needs of environments with limited space, budget and/or IT personnel constraints.



## 5.0 References

- vSAN 2-Node Guide (<https://storagehub.vmware.com/t/vmware-vsan/vsan-2-node-guide/>)
- vSAN Stretched Cluster (<https://storagehub.vmware.com/t/vmware-vsan/vsan-stretched-cluster-guide/>)
- 2-NODE VSAN – WITNESS NETWORK DESIGN CONSIDERATIONS (<https://cormachogan.com/2017/10/06/2-node-vsan-witness-network-design-considerations/>)
- vSAN Stretched Cluster Bandwidth Sizing (<https://storagehub.vmware.com/t/vmware-vsan/vsan-stretched-cluster-bandwidth-sizing/>)
- VxRail Network Guide (<https://www.dell.com/resources/en-us/asset/technical-guides-support-information/products/converged-infrastructure/h15300-vxrail-network-guide.pdf>)
- VxRail vCenter Server Planning Guide (<https://www.dell.com/resources/en-us/asset/technical-guides-support-information/products/converged-infrastructure/vxrail-vcenter-server-planning-guide.pdf>)



# Processador Intel® Xeon® Gold 6230

cache de 27,5 M, 2,10 GHz

## Especificações

Especificações de exportação

### Essenciais

Coleção de produtos	Processadores escaláveis Intel® Xeon® da 2ª Geração
Codiname	Produtos com denominação anterior Cascade Lake
Segmento vertical	Server
Número do processador	6230
Status	Launched
Data de introdução	Q2'19
Litografia	14 nm
Condições de uso	Server/Enterprise
Preço recomendado para o cliente	\$1894.00 - \$1900.00

### Desempenho

Número de núcleos	20
Nº de threads	40
Frequência baseada em processador	2.10 GHz
Frequência turbo max	3.90 GHz
Cache	27.5 MB
Nº de links de UPI	3



TDP

Brasil (Português)

125W



## Informações complementares

Opções integradas disponíveis

Sim

## Especificações de memória

Tamanho máximo de memória (de acordo com o tipo de memória)

1 TB

Tipos de memória

DDR4-2933

Velocidade máxima de memória

2933 MHz

Nº máximo de canais de memória

6

Compatibilidade com memória ECC †

Sim

Memória persistente Intel® Optane™ DC com suporte

Sim

## Opções de expansão

Escalabilidade

4S

Revisão de PCI Express

3.0

Nº máximo de linhas PCI Express

48

## Especificações de encapsulamento

Soquetes suportados

FCLGA3647

T<sub>CASE</sub>

87°C

Tamanho do pacote

76.0mm x 56.5mm

## Tecnologias avançadas

Intel® Deep Learning Boost (Intel® DL Boost)





Sim

Intel® Speed Select Technology – Perfil de desempenho

Não

Intel® Speed Select Technology – Frequência básica

Não

 Intel® Resource Director Technology (Intel® RDT) <span style="float: right;">             Brasil (Português)   Sim              </span>	
Tecnologia Intel® Speed Shift	Sim
Tecnologia Intel® Turbo Boost Max 3.0 ‡	Não
Tecnologia Intel® Turbo Boost ‡	2.0
Elegibilidade da plataforma Intel® vPro™ ‡	Sim
Tecnologia Hyper-Threading Intel® ‡	Sim
Tecnologia de virtualização Intel® (VT-x) ‡	Sim
Tecnologia de virtualização Intel® para E/S direcionada (VT-d) ‡	Sim
Intel® VT-x com Tabelas de página estendida (EPT) ‡	Sim
Intel® TSX-NI	Sim
Intel® 64 ‡	Sim
Extensões do conjunto de instruções	Intel® SSE4.2, Intel® AVX, Intel® AVX2, Intel® AVX-512
Nº de unidades de FMA de AVX-512	2
Tecnologia Enhanced Intel SpeedStep®	Sim
Intel® Volume Management Device (VMD - Dispositivo de Gerenciamento de Volume)	Sim

## Segurança e confiabilidade

Novas instruções Intel® AES	Sim
Intel® Trusted Execution Technology ‡	Sim
Bit de desativação de execução ‡	Sim
Tecnologia Intel® Run Sure	Sim
Controle de Execução baseado em Modo (MBE — Mode-based Execute Control)	Sim

Pedidos e conformidade

Produtos compatíveis

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.

As classificações da Intel são apenas para fins informativos e consistem em Export Control Classification Numbers (ECCN — Número de Classificação de Controle de Exportações) e Harmonized Tariff Schedule (HTS — Programa de Tarifas Harmonizadas). Quaisquer usos das classificações da Intel são sem os recursos da Intel e não devem ser interpretados como uma representação ou garantia relacionada ao ECCN ou HTS apropriado. Como exportadora e/ou importadora, sua empresa é responsável por determinar a classificação correta de sua transação

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.

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.

SKUs "anunciados" ainda não estão disponíveis. Favor consultar a data de lançamento para a disponibilidade no mercado.

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.

Os processadores compatíveis com a computação de 64 bits na arquitetura Intel® requerem BIOS habilitados para arquitetura Intel 64.

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 [www.intel.com/content/www/br/pt/architecture-and-technology/turbo-boost/turbo-boost-technology.html](http://www.intel.com/content/www/br/pt/architecture-and-technology/turbo-boost/turbo-boost-technology.html)

Consulte <http://www.intel.com/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 números dos processadores Intel não são indicação de desempenho. Os valores do processador diferenciam recursos dentro de cada família de processador e não entre famílias diferentes de processador. Consulte <http://www.intel.com.br/content/www/br/pt/processors/processor-numbers.html>

para obter mais detalhes.


Informações sobre a empresa

Nosso compromisso

Relações com investidores

Fale conosco

Sala de imprensa

 Mapa do site

Brasil (Português) 



Empregos



© Intel Corporation

[Termos de uso](#)

[\\*Marcas comerciais](#)

[Privacidade](#)

[Cookies](#)

[Transparência da cadeira de fornecimento](#)

# vSphere Monitoring and Performance

17 APR 2018

VMware vSphere 6.7

VMware ESXi 6.7

vCenter Server 6.7



vmware®



You can find the most up-to-date technical documentation on the VMware website at:

<https://docs.vmware.com/>

If you have comments about this documentation, submit your feedback to

[docfeedback@vmware.com](mailto:docfeedback@vmware.com)

**VMware, Inc.**  
3401 Hillview Ave.  
Palo Alto, CA 94304  
[www.vmware.com](http://www.vmware.com)

Copyright © 2010–2018 VMware, Inc. All rights reserved. [Copyright and trademark information.](#)

# Contents

About vSphere Monitoring and Performance	5
<b>1 Monitoring Inventory Objects with Performance Charts</b>	<b>7</b>
Performance Chart Types	8
Data Counters	8
Metric Groups in vSphere	10
Data Collection Intervals	11
Data Collection Levels	11
View Performance Charts	12
Performance Charts Options Available Under the View Menu	13
Overview Performance Charts	14
Working with Advanced and Custom Charts	106
Troubleshoot and Enhance Performance	109
<b>2 Monitoring Guest Operating System Performance</b>	<b>116</b>
Enable Statistics Collection for Guest Operating System Performance Analysis	116
View Performance Statistics for Windows Guest Operating Systems	116
<b>3 Monitoring Host Health Status</b>	<b>118</b>
Monitor Health Status in the vSphere Client	119
Monitor Hardware Health Status	119
Reset Health Status Sensors in the vSphere Client	120
Reset Health Status Sensors	120
<b>4 Monitoring vSphere Health</b>	<b>121</b>
Check vSphere Health in vSphere Client	121
<b>5 Monitoring Events, Alarms, and Automated Actions</b>	<b>123</b>
View Events	125
View System Logs	125
Export Events Data	126
Streaming Events to a Remote Syslog Server	126
Retention of Events in the vCenter Server Database	128
View Triggered Alarms and Alarm Definitions	130
Live Refresh of Recent Tasks and Alarms	130
Set an Alarm in the vSphere Web Client	131
Set an Alarm in the vSphere Client	141
Acknowledge Triggered Alarms	144

- Reset Triggered Event Alarms 144
- Preconfigured vSphere Alarms 145
  
- 6 Monitoring Solutions with the vCenter Solutions Manager 151**
  - View Solutions and vServices 151
  - Monitoring Agents 152
  - Monitoring vServices 152
  
- 7 Monitoring the Health of Services and Nodes 154**
  - View the Health Status of Services and Nodes 154
  
- 8 Performance Monitoring Utilities: resxtop and esxtop 156**
  - Using the esxtop Utility 156
  - Using the resxtop Utility 157
  - Using esxtop or resxtop in Interactive Mode 158
  - Using Batch Mode 173
  - Using Replay Mode 174
  
- 9 Using the vimtop Plug-In to Monitor the Resource Use of Services 177**
  - Monitor Services By Using vimtop in Interactive Mode 177
  - Interactive Mode Command-Line Options 177
  - Interactive Mode Single-Key Commands for vimtop 178
  
- 10 Monitoring Networked Devices with SNMP and vSphere 180**
  - Using SNMP Traps with vCenter Server 180
  - Configure SNMP for ESXi 182
  - SNMP Diagnostics 193
  - Monitor Guest Operating Systems with SNMP 193
  - VMware MIB Files 193
  - SNMPv2 Diagnostic Counters 195
  
- 11 System Log Files 197**
  - View System Log Entries 197
  - View System Logs on an ESXi Host 197
  - System Logs 198
  - Export System Log Files 199
  - ESXi Log Files 200
  - Upload Logs Package to a VMware Service Request 200
  - Configure Syslog on ESXi Hosts 201
  - Configuring Logging Levels for the Guest Operating System 202
  - Collecting Log Files 204
  - Viewing Log Files with the Log Browser 208

# About vSphere Monitoring and Performance

VMware provides several tools to help you monitor your virtual environment and to locate the source of potential issues and current problems.

<b>Performance charts</b>	Allow you to see performance data on a variety of system resources including CPU, Memory, Storage, and so on.
<b>Performance monitoring command-line utilities</b>	Allow you to access detailed information on system performance through the command line.
<b>Host health</b>	Allows you to quickly identify which hosts are healthy and which are experiencing problems.
<b>Events, alerts, and alarms</b>	Allow you to configure alerts and alarms and to specify the actions the system should take when they are triggered.
<b>System Log Files</b>	System logs contain additional information about activities in your vSphere environment.

## Intended Audience

The content in this section is intended for vSphere administrators who perform the following tasks:

- Monitor the health and performance of physical hardware backings for the virtual environment.
- Monitor the health and performance of virtual devices in the virtual environment.
- Troubleshoot problems in the system.
- Configure alarms.
- Configure SNMP messages.

Virtual machine administrators also might find the section on [Chapter 2 Monitoring Guest Operating System Performance](#) helpful.

## vSphere Web Client and vSphere Client

Instructions in this guide reflect the vSphere Client (an HTML5-based GUI). You can also use the instructions to perform most of the tasks by using the vSphere Web Client (a Flex-based GUI).

Tasks for which the workflow differs significantly between the vSphere Client and the vSphere Web Client have duplicate procedures that provide steps according to the respective client interface. The procedures that relate to the vSphere Web Client, contain vSphere Web Client in the title.

---

**Note** In vSphere 6.7, most of the vSphere Web Client functionality is implemented in the vSphere Client. For an up-to-date list of the unsupported functionality, see [Functionality Updates for the vSphere Client](#).

---

## VMware Technical Publications Glossary

VMware Technical Publications provides a glossary of terms that might be unfamiliar to you. For definitions of terms as they are used in VMware technical documentation, go to

<http://www.vmware.com/support/pubs>.

# Monitoring Inventory Objects with Performance Charts



The vSphere statistics subsystem collects data on the resource usage of inventory objects. Data on a wide range of metrics is collected at frequent intervals, processed, and archived in the vCenter Server database. You can access statistical information through command-line monitoring utilities or by viewing performance charts in the vSphere Web Client.

## Counters and Metric Groups

vCenter Server systems and hosts use data counters to query for statistics. A data counter is a unit of information relevant to a given inventory object or device. Each counter collects data for a different statistic in a metric group. For example, the disk metric group includes separate data counters to collect data for disk read rate, disk write rate, and disk usage. Statistics for each counter are rolled up after a specified collection interval. Each data counter consists of several attributes that are used to determine the statistical value collected.

For a complete list and description of performance metrics, see the *vSphere API Reference*.

---

**Note** Counters that are introduced in later versions might not contain data from hosts of earlier versions. For details, see the VMware Knowledge Base.

---

## Collection Levels and Collection Intervals

Collection levels determine the number of counters for which data is gathered during each collection interval. Collection intervals determine the time period during which statistics are aggregated, calculated, rolled up, and archived in the vCenter Server database. Together, the collection interval and collection level determine how much statistical data is collected and stored in your vCenter Server database.

## Data Availability

Real-time data appears in the performance charts only for hosts and virtual machines that are powered on. Historical data appears for all supported inventory objects, but might be unavailable during certain circumstances.

This chapter includes the following topics:

- [Performance Chart Types](#)
- [Data Counters](#)
- [Metric Groups in vSphere](#)
- [Data Collection Intervals](#)
- [Data Collection Levels](#)
- [View Performance Charts](#)
- [Performance Charts Options Available Under the View Menu](#)
- [Overview Performance Charts](#)
- [Working with Advanced and Custom Charts](#)
- [Troubleshoot and Enhance Performance](#)

## Performance Chart Types

Performance metrics are displayed in different types of charts, depending on the metric type and object.

**Table 1-1. Performance Chart Types**

Chart Type	Description
Line chart	Displays metrics for a single inventory object. The data for each performance counter is plotted on a separate line in the chart. For example, a network chart for a host can contain two lines: one showing the number of packets received, and one showing the number of packets transmitted.
Bar chart	Displays storage metrics for datastores in a selected data center. Each datastore is represented as a bar in the chart. Each bar displays metrics based on the file type: virtual disks, snapshots, swap files, and other files.
Pie chart	Displays storage metrics for a single object, based on the file types, or virtual machines. For example, a pie chart for a datastore can display the amount of storage space occupied by the virtual machines taking up the largest space.
Stacked chart	Displays metrics for the child objects that have the highest statistical values. All other objects are aggregated, and the sum value is displayed with the term <b>Other</b> . For example, a host's stacked CPU usage chart displays CPU usage metrics for the 10 virtual machines on the host that are consuming the most CPU. The <b>Other</b> amount contains the total CPU usage of the remaining virtual machines.  The metrics for the host itself are displayed in separate line charts.  Stacked charts are useful in comparing the resource allocation and usage across multiple hosts or virtual machines. By default, the 10 child objects with the highest data counter values are displayed.

## Data Counters

Each data counter includes several attributes that are used to determine the statistical value collected. See the *vSphere API Reference* for a complete list and description of supported counters.

**Table 1-2. Data Counter Attributes**

Attribute	Description
Unit of Measurement	<p>Standard in which the statistic quantity is measured.</p> <ul style="list-style-type: none"> <li>■ Kilobytes (KB) – 1024 bytes</li> <li>■ Kilobytes per second (KBps) – 1024 bytes per second</li> <li>■ Kilobits (kb) – 1000 bits</li> <li>■ Kilobits per second (kbps) – 1000 bits per second</li> <li>■ Megabytes (MB)</li> <li>■ Megabytes per second (MBps)</li> <li>■ Megabits (Mb), megabits per second (Mbps)</li> <li>■ Megahertz (MHz)</li> <li>■ Microseconds (µs)</li> <li>■ Milliseconds (ms)</li> <li>■ Number (#)</li> <li>■ Percent (%)</li> <li>■ Seconds (s)</li> </ul>
Description	Text description of the data counter.
Statistics Type	<p>Measurement used during the statistics interval. Related to the unit of measurement.</p> <ul style="list-style-type: none"> <li>■ Rate – Value over the current statistics interval</li> <li>■ Delta – Change from previous statistics interval.</li> <li>■ Absolute – Absolute value (independent of the statistics interval).</li> </ul>
Rollup Type	<p>Calculation method used during the statistics interval to roll up data. Determines the type of statistical values that are returned for the counter.</p> <ul style="list-style-type: none"> <li>■ Average – Data collected during the interval is aggregated and averaged.</li> <li>■ Minimum – The minimum value is rolled up.</li> <li>■ Maximum – The maximum value is rolled up.</li> </ul> <p>The Minimum and Maximum values are collected and displayed only in statistics level 4. Minimum and maximum rollup types are used to capture peaks in data during the interval. For real-time data, the value is the current minimum or current maximum. For historical data, the value is the average minimum or average maximum.</p> <p>For example, the following information for the CPU usage chart shows that the average is collected at statistics level 1. The minimum and maximum values are collected at statistics level 4.</p> <ul style="list-style-type: none"> <li>■ Counter: usage</li> <li>■ Unit: Percentage (%)</li> <li>■ Rollup Type: Average (Minimum/Maximum)</li> <li>■ Collection Level: 1 (4)</li> </ul> <ul style="list-style-type: none"> <li>■ Summation – Data collected is summed. The measurement displayed in the chart represents the sum of data collected during the interval.</li> <li>■ Latest – Data collected during the interval is a set value. The value displayed in the performance charts represents the current value.</li> </ul>
Collection level	<p>Number of data counters used to collect statistics. Collection levels range from 1 to 4, with 4 having the most counters.</p> <p><b>Note</b> Be careful when you set a higher collection level, as the process requires significant increase of resource usage. For more information, see <a href="#">Data Collection Levels</a>.</p>



## Metric Groups in vSphere

The performance data collection subsystem for vSphere collects performance data on various inventory items and their devices. Data counters define individual performance metrics. Performance metrics are organized into logical groups based on the object or object device. Statistics for one or more metrics can be displayed in a chart.

**Table 1-3. Metric Groups**

Metric group	Description
Cluster Services	Performance statistics for clusters configured by using vSphere Distributed Resource Scheduler, vSphere High Availability, or both.
CPU	CPU utilization per host, virtual machine, resource pool, or compute resource.
Datastore	Statistics for datastore utilization.
Disk	Disk utilization per host, virtual machine, or datastore. Disk metrics include I/O performance, such as latency and read/write speeds, and utilization metrics for storage as a finite resource.
Memory	Memory utilization per host, virtual machine, resource pool, or compute resource. The value obtained is one of the following: <ul style="list-style-type: none"> <li>■ For virtual machines, memory refers to the guest physical memory. Guest physical memory is the amount of physical memory presented as a virtual-hardware component to the virtual machine, at creation time, and made available when the virtual machine is running.</li> <li>■ For hosts, memory refers to the machine memory. Machine memory is the RAM that is installed on the hardware that comprises the host.</li> </ul>
Network	Network utilization for both physical and virtual network interface controllers (NICs) and other network devices. The virtual switches that support connectivity among all components, such as hosts, virtual machines, VMkernel.
Power	Energy usage statistics per host.
Storage Adapter	Data traffic statistics per host bus adapter (HBA).
Storage Path	Data traffic statistics per path.
System	Overall system availability, such as the system heartbeat and uptime. These counters are available directly from hosts and from vCenter Server.
Virtual Disk	Disk utilization and disk performance metrics for virtual machines.
Virtual Flash	Virtual flash counters.
Virtual Machine Operations	Virtual machine power and provisioning operations in a cluster or data center.
vSphere Replication	Statistics for the virtual machine replication performed by VMware vCenter Site Recovery Manager.

## Data Collection Intervals

Collection intervals determine the duration for which statistics are aggregated, calculated, rolled up, and archived. Together, the collection interval and collection level determine how much statistical data is gathered and stored in your vCenter Server database.

**Table 1-4. Collection Intervals**

Collection Interval/Archive Length	Collection Frequency	Default Behavior
1 Day	5 Minutes	<p>Real-time statistics are rolled up to create one data point every 5 minutes. The result is 12 data points every hour and 288 data points every day. After 30 minutes, the six data points collected are aggregated and rolled up as a data point for the 1-Week time range.</p> <p>You can change the interval duration and archive length of the 1-Day collection interval by configuring the statistics settings.</p>
1 Week	30 Minutes	<p>1-Day statistics are rolled up to create one data point every 30 minutes. The result is 48 data points every day and 336 data points every week. Every 2 hours, the 12 data points collected are aggregated and rolled up as a data point for the 1-Month time range.</p> <p>You cannot change the default settings of the 1-Week collection interval.</p>
1 Month	2 Hours	<p>1-Week statistics are rolled up to create one data point every 2 hours. The result is 12 data points every day and 360 data points every month (assuming a 30-day month). After 24 hours, the 12 data points collected are aggregated and rolled up as a data point for the <b>1-Year</b> time range.</p> <p>You cannot change the default settings of the 1-Month collection interval.</p>
1 Year	1 Day	<p>1-Month statistics are rolled up to create one data point every day. The result is 365 data points each year.</p> <p>You can change the archive length of the 1-Year collection interval by configuring the statistics settings.</p>

---

**Note** If you change the duration of data collection intervals you might need to allocate more storage resources.

---

## Data Collection Levels

Each collection interval has a default collection level that determines the amount of data gathered and which counters are available for display in the charts. Collection levels are also referred to as statistics levels.

**Table 1-5. Statistics Levels**

Level	Metrics	Best Practice
Level 1	<ul style="list-style-type: none"> <li>■ Cluster Services (VMware Distributed Resource Scheduler) – all metrics</li> <li>■ CPU – cpuentitlement, totalmhz, usage (average), usagemhz</li> <li>■ Disk – capacity, maxTotalLatency, provisioned, unshared, usage (average), used</li> <li>■ Memory – consumed, mementitlement, overhead, swapinRate, swapoutRate, swapused, totalmb, usage (average), vmmemctl (balloon)</li> <li>■ Network – usage (average), IPv6</li> <li>■ System – heartbeat, uptime</li> <li>■ Virtual Machine Operations – numChangeDS, numChangeHost, numChangeHostDS</li> </ul>	<p>Use for long-term performance monitoring when device statistics are not required.</p> <p>Level 1 is the default Collection Level for all Collection Intervals.</p>
Level 2	<ul style="list-style-type: none"> <li>■ Level 1 metrics</li> <li>■ CPU – idle, reservedCapacity</li> <li>■ Disk – All metrics, excluding numberRead and numberWrite.</li> <li>■ Memory – All metrics, excluding memUsed and maximum and minimum rollup values.</li> <li>■ Virtual Machine Operations – All metrics</li> </ul>	<p>Use for long-term performance monitoring when device statistics are not required but you want to monitor more than the basic statistics.</p>
Level 3	<ul style="list-style-type: none"> <li>■ Level 1 and Level 2 metrics</li> <li>■ Metrics for all counters, excluding minimum and maximum rollup values.</li> <li>■ Device metrics</li> </ul>	<p>Use for short-term performance monitoring after encountering problems or when device statistics are required.</p> <p>Because of the large quantity of troubleshooting data retrieved and recorded, use level 3 for the shortest time period (Day or Week collection interval).</p>
Level 4	All metrics supported by the vCenter Server, including minimum and maximum rollup values.	<p>Use for short-term performance monitoring after encountering problems or when device statistics are required.</p> <p>Because of the large quantity of troubleshooting data retrieved and recorded, use level 4 for the shortest amount of time.</p>

---

**Note** When you increase the collection level, the storage and system requirements might change. You might need to allocate more system resources to avoid a decrease in the performance.

---

## View Performance Charts

The vCenter Server statistics settings, the type of object selected, and the features that are enabled on the selected object determine the amount of information displayed in charts. Charts are organized into views. You can select a view to see related data together on one screen. You can also specify the time range, or data collection interval. The duration extends from the selected time range to the present time.

Overview charts display multiple data sets in one panel to evaluate different resource statistics, display thumbnail charts for child objects. It also displays charts for a parent and a child object. Advanced charts display more information than overview charts, are configurable, and can be printed or exported. You can export data in the PNG, JPEG, or CSV formats. See [View Advanced Performance Charts](#).

### Procedure

- 1 Select a valid inventory object in the vSphere Web Client.

Overview and advanced performance charts are available for datacenter, cluster, host, resource pool, vApp, and virtual machine objects. Overview charts are also available for datastores and datastore clusters. Performance charts are not available for network objects.

- 2 Click the **Monitor** tab, and click **Performance**.

- 3 Select a view.

Available views depend on the type of object. For views that might contain many charts in a large environment, the vSphere Web Client displays the charts distributed on multiple pages. You can use the arrow buttons to navigate between pages.

- 4 Select a predefined or custom time range.

## Performance Charts Options Available Under the View Menu

The performance chart options that you can access under the **View** menu vary depending on the type of inventory object you select.

For example, the **Virtual Machines** view is available when you view host performance charts only if there are virtual machines on the selected host. Likewise, the **Fault Tolerance** view for virtual machine performance charts is available only when that feature is enabled for the selected virtual machine.

**Table 1-6. Performance Chart Views by Inventory Object**

Object	View List Items
Data center	<ul style="list-style-type: none"> <li>■ <b>Storage</b> - space utilization charts for datastores in the data center, including space by file type and storage space used by each datastore in the data center.</li> <li>■ <b>Clusters</b> - thumbnail CPU and memory charts for each cluster, and stacked charts for total CPU and memory usage in the data center. This view is the default.</li> </ul>
Datastore and datastore cluster	<ul style="list-style-type: none"> <li>■ <b>Space</b> - space utilization charts for the datastore: <ul style="list-style-type: none"> <li>■ space utilization by file type</li> <li>■ space utilization by virtual machine</li> <li>■ space usage</li> </ul> </li> <li>■ <b>Performance</b> - performance charts for the datastore or datastore cluster and for virtual machine disks on the resource.</li> </ul>

**Note** The Performance view for datastores is only available when all hosts that are connected to the datastores are ESX/ESXi 4.1 or greater. The Performance view for datastore clusters is only available when the Storage DRS is enabled.

**Table 1-6. Performance Chart Views by Inventory Object (Continued)**

Object	View List Items
Cluster	<ul style="list-style-type: none"> <li>■ <b>Home</b> - CPU and memory charts for the cluster.</li> <li>■ <b>Resource Pools &amp; Virtual Machines</b> - thumbnail charts for resource pools and virtual machines, and stacked charts for total CPU and memory usage in the cluster.</li> <li>■ <b>Hosts</b> - thumbnail charts for each host in the cluster, and stacked charts for total CPU, memory, disk usage, and network usage.</li> </ul>
Host	<ul style="list-style-type: none"> <li>■ <b>Home</b> - CPU, memory, disk, and network charts for the host.</li> <li>■ <b>Virtual Machines</b> - thumbnail charts for virtual machines, and stacked charts for total CPU usage and total memory usage on the host.</li> </ul>
Resource Pool and vApps	<ul style="list-style-type: none"> <li>■ <b>Home</b> - CPU and memory charts for the resource pool.</li> <li>■ <b>Resource Pools &amp; Virtual Machines</b> - thumbnail charts for resource pools, and virtual machines and stacked charts for CPU and memory usage in the resource pool or vApp.</li> </ul>
Virtual Machine	<ul style="list-style-type: none"> <li>■ <b>Storage</b> - space utilization charts for the virtual machine: space by file type, space by datastore, and total gigabytes.</li> <li>■ <b>Fault Tolerance</b> - CPU and memory charts that display comparative metrics for the fault-tolerant primary and secondary virtual machines.</li> <li>■ <b>Home</b> - CPU, memory, network, host (thumbnail charts), and disk usage charts for the virtual machine.</li> </ul>

## Overview Performance Charts

The overview performance charts display the most common metrics for an object in the inventory. Use these charts to monitor and troubleshoot performance problems.

The metrics provided in Overview performance charts are a subset of those collected for hosts and the vCenter Server. For a complete list of all metrics collected by hosts and the vCenter Server, see the *vSphere API Reference*.

## Clusters

The cluster charts contain information about CPU, disk, memory, and network usage for clusters. The help topic for each chart contains information about the data counters displayed in that chart. The collection level set for vCenter Server determines the available counters.

### CPU (MHz)

The CPU (MHz) chart displays CPU usage for the cluster.

#### Cluster Counters

This chart is located in the Home view of the Cluster **Performance** tab.

**Table 1-7. Data Counters**

Chart Label	Description
Usage	<p>Sum of the average CPU usage values, in Megahertz, of all virtual machines in the cluster.</p> <ul style="list-style-type: none"> <li>■ Counter: usagemhz</li> <li>■ Stats Type: Rate</li> <li>■ Unit: Megahertz (MHz)</li> <li>■ Rollup Type: Average (Minimum/Maximum)</li> <li>■ Collection Level: 1 (4)</li> </ul>
Total	<p>Total amount of CPU resources available in the cluster. The maximum value is equal to the number of cores multiplied by the frequency of the processors.</p> <p>For example, a cluster has two hosts, each of which has four CPUs that are 3 GHz each, and one virtual machine that has two virtual CPUs.</p> <p>VM totalmhz = 2 vCPUs * 3000 MHz = 6000 MHz</p> <p>Host totalmhz = 4 CPUs * 3000 MHz = 12000 MHz</p> <p>Cluster totalmhz = 2 x 4 * 3000 MHz = 24000 MHz</p> <ul style="list-style-type: none"> <li>■ Counter: totalmhz</li> <li>■ Stats Type: Rate</li> <li>■ Unit: Megahertz (MHz)</li> <li>■ Rollup Type: Summation</li> <li>■ Collection Level: 1</li> </ul>

### Chart Analysis

A short spike in CPU usage indicates that you are making the best use of cluster resources. However, if the value is constantly high, the CPU demanded is likely greater than the CPU capacity available. A high CPU usage value can lead to increased ready time and processor queuing of the virtual machines on the hosts in the cluster.

If performance is impacted, consider taking the following actions.

**Table 1-8. CPU Performance Enhancement Advice**

#	Resolution
1	Verify that VMware Tools is installed on each virtual machine.
2	<p>If the cluster is not a DRS cluster, enable DRS. To enable DRS, perform the following tasks:</p> <ol style="list-style-type: none"> <li>1 Select the cluster, and click the <b>Configure</b> tab.</li> <li>2 Under <b>Services</b>, click <b>vSphere DRS</b>.</li> <li>3 click <b>Edit</b>.</li> </ol> <p>An Edit Cluster Settings dialog box opens.</p> <ol style="list-style-type: none"> <li>4 Click <b>Turn ON vSphere DRS</b>, and click <b>OK</b>.</li> </ol>
3	<p>If the cluster is a DRS cluster:</p> <ul style="list-style-type: none"> <li>■ Increase the number of hosts, and migrate one or more virtual machines to the new host.</li> <li>■ Check the aggressiveness threshold. If the value is low, increase the threshold. This might help avoid hot spots in the cluster.</li> </ul>
4	Migrate one or more virtual machines to a new host.
5	Upgrade the physical CPUs or cores on each host in the cluster if necessary.

**Table 1-8. CPU Performance Enhancement Advice (Continued)**

#	Resolution
6	Enable CPU-saving features, such as TCP Segmentation Offload.
7	Replace software I/O with the dedicated hardware, such as iSCSI HBAs or TCP Segmentation Offload NICs.

## CPU Usage

The cluster CPU Usage charts monitor the CPU utilization of the hosts, resource pools, and virtual machines in the cluster. This chart displays the 10 child objects in the cluster with the most CPU usage.

This chart is located in the Resource Pools and Virtual Machines view of the Cluster **Performance** tab.

**Table 1-9. Data Counters**

Chart Label	Description
<host>, <resource pool>, or <virtual machine>	<p>Amount of CPU actively used by the host, resource pool, or virtual machine in the cluster.</p> <ul style="list-style-type: none"> <li>■ Counter: usagemhz</li> <li>■ Stats Type: Rate</li> <li>■ Unit: MegaHertz (MHz)</li> <li>■ Rollup Type: Average (Minimum/Maximum)</li> <li>■ Collection Level: 1 (4)</li> </ul>

## Chart Analysis

A short spike in CPU usage indicates that you are making the best use of cluster resources. However, if the value is constantly high, the CPU demanded is likely greater than the CPU capacity available. A high CPU usage value can lead to increased ready time and processor queuing of the virtual machines on the hosts in the cluster.

If performance is impacted, consider taking the following actions.

**Table 1-10. CPU Performance Enhancement Advice**

#	Resolution
1	Verify that VMware Tools is installed on each virtual machine.
2	<p>If the cluster is not a DRS cluster, enable DRS. To enable DRS, perform the following tasks:</p> <ol style="list-style-type: none"> <li>1 Select the cluster, and click the <b>Configure</b> tab.</li> <li>2 Under <b>Services</b>, click <b>vSphere DRS</b>.</li> <li>3 click <b>Edit</b>.</li> </ol> <p>An Edit Cluster Settings dialog box opens.</p> <ol style="list-style-type: none"> <li>4 Click <b>Turn ON vSphere DRS</b>, and click <b>OK</b>.</li> </ol>
3	<p>If the cluster is a DRS cluster:</p> <ul style="list-style-type: none"> <li>■ Increase the number of hosts, and migrate one or more virtual machines to the new host.</li> <li>■ Check the aggressiveness threshold. If the value is low, increase the threshold. This might help avoid hot spots in the cluster.</li> </ul>
4	Migrate one or more virtual machines to a new host.
5	Upgrade the physical CPUs or cores on each host in the cluster if necessary.

**Table 1-10. CPU Performance Enhancement Advice (Continued)**

#	Resolution
6	Enable CPU-saving features, such as TCP Segmentation Offload.
7	Replace software I/O with the dedicated hardware, such as iSCSI HBAs or TCP Segmentation Offload NICs.

## Disk (KBps)

The Disk (KBps) chart displays the disk I/O of the 10 hosts in the cluster with the most disk usage.

This chart is located in the Hosts view of the cluster **Performance** tab.

**Table 1-11. Data Counters**

Chart Label	Description
<i>host_name</i>	<p>Average data I/O rate across all hosts in the cluster.</p> <ul style="list-style-type: none"> <li>■ Counter: usage</li> <li>■ Stats Type: Rate</li> <li>■ Unit: Kilobytes per second (KBps)</li> <li>■ Rollup Type: Average (Minimum/Maximum)</li> <li>■ Collection Level: 1 (4)</li> </ul>

## Chart Analysis

Use the disk charts to monitor average disk loads and to determine trends in disk usage. For example, you might notice a performance degradation with applications that frequently read from and write to the hard disk. If you see a spike in the number of disk read or write requests, check whether any such applications were running then.

The best ways to determine if your vSphere environment is experiencing disk problems is to monitor the disk latency data counters. You can use the advanced performance charts to view these statistics.

- The `kernelLatency` data counter measures the average amount of time, in milliseconds, that the VMkernel spends processing each SCSI command. For best performance, the value must be 0 -1 milliseconds. If the value is greater than 4 ms, the virtual machines on the host are trying to send more throughput to the storage system than the configuration supports. Check the CPU usage, and increase the queue depth.
- The `deviceLatency` data counter measures the average amount of time, in milliseconds, to complete a SCSI command from the physical device. Depending on your hardware, a number greater than 15 ms indicates probable problems with the storage array. Move the active VMDK to a volume with more spindles or add disks to the LUN.
- The `queueLatency` data counter measures the average amount of time taken per SCSI command in the VMkernel queue. This value must always be zero. If not, the workload is too high and the array cannot process the data fast enough.

If the disk latency values are high, or if you notice other problems with disk I/O performance, consider taking the following actions.



**Table 1-12. Disk I/O Performance Enhancement Advice**

#	Resolution
1	<p>Increase the virtual machine memory. It allows more operating system caching, which reduces I/O activity. Note: It might require you to increase the host memory. Increasing memory might reduce the need to store data because databases can utilize the system memory to cache data and avoid disk access.</p> <p>To verify that virtual machines have adequate memory, check swap statistics in the guest operating system. Increase the guest memory, but not to an extent that leads to excessive host memory swapping. Install VMware Tools so that memory ballooning can occur.</p>
2	Defragment the file systems on all guests.
3	Disable antivirus on-demand scans on the VMDK and VMEM files.
4	Use the vendor's array tools to determine the array performance statistics. When too many servers simultaneously access common elements on an array, the disks might have trouble keeping up. To increase throughput, consider array-side improvements.
5	Use Storage vMotion to migrate I/O-intensive virtual machines across multiple hosts.
6	Balance the disk load across all physical resources available. Spread heavily used storage across LUNs that are accessed by different adapters. Use separate queues for each adapter to improve disk efficiency.
7	Configure the HBAs and RAID controllers for optimal use. Verify that the queue depths and cache settings on the RAID controllers are adequate. If not, increase the number of outstanding disk requests for the virtual machine by adjusting the <code>Disk.SchedNumReqOutstanding</code> parameter. For more information, see <i>vSphere Storage</i> .
8	For resource-intensive virtual machines, separate the virtual machine's physical disk drive from the drive with the system page file. It alleviates disk spindle contention during periods of high use.
9	On systems with sizable RAM, disable memory trimming by adding the line <code>MemTrimRate=0</code> to the virtual machine's VMX file.
10	If the combined disk I/O is higher than a single HBA capacity, use multipathing or multiple links.
11	For ESXi hosts, create virtual disks as preallocated. When you create a virtual disk for a guest operating system, select <b>Allocate all disk space now</b> . The performance degradation associated with reassigning additional disk space does not occur, and the disk is less likely to become fragmented.
12	Use the most current host hardware.

## Memory (MB)

The Memory (MB) chart displays consumed memory for the cluster. The chart appears only at collection level 1.

This chart is located in the Home view of the cluster **Performance** tab.

**Table 1-13. Data Counters**

Chart Label	Description
Consumed	<p>Amount of host machine memory used by all powered on virtual machines in the cluster. A cluster's consumed memory consists of virtual machine consumed memory and overhead memory. It does not include host-specific overhead memory, such as memory used by the service console or VMkernel.</p> <ul style="list-style-type: none"> <li>■ Counter: consumed</li> <li>■ Stats Type: Absolute</li> <li>■ Unit: Megabytes (MB)</li> <li>■ Rollup Type: Average (Minimum/Maximum)</li> <li>■ Collection Level: 1 (4)</li> </ul>
Total	<p>Total amount of machine memory of all hosts in the cluster that is available for virtual machine memory (physical memory for use by the Guest OS) and virtual machine overhead memory.</p> <p>Memory Total = Aggregate host machine memory - (VMkernel memory + Service Console memory + other service memory)</p> <p><b>Note</b> The totalmb data counter is the same as the effectivemem data counter, which is supported only for backward compatibility.</p> <ul style="list-style-type: none"> <li>■ Counter: totalmb</li> <li>■ Stats Type: Absolute</li> <li>■ Unit: Megabytes (MB)</li> <li>■ Rollup Type: Average (Minimum/Maximum)</li> <li>■ Collection Level: 1 (4)</li> </ul>

### Chart Analysis

Memory usage is not an indicator of performance problems. Memory can be high if a host is swapping or ballooning, which can result in virtual machine guest swapping. In such cases, check for other problems, such as CPU over-commitment or storage latencies.

If you have constantly high memory usage in a cluster, resource pool, or vApp, consider taking the following actions.

**Table 1-14. Memory Performance Enhancement Advice**

#	Resolution
1	Verify that VMware Tools is installed on each virtual machine.
2	Verify that the balloon driver is enabled. The balloon driver is installed with VMware Tools and is critical to performance. The VMkernel regularly reclaims unused virtual machine memory by ballooning and swapping. Generally, it does not impact virtual machine performance.
3	If the balloon value is high, check the resource shares, reservations, and limits for the virtual machines and resource pools on the hosts. Verify that the host's settings are adequate and not lower than those set for the virtual machine. If free memory is available on the hosts and the virtual machines are experiencing high swap or balloon memory, the virtual machine (or resource pool, if it belongs to one) has reached its resource limit. Check the maximum resource limit set on that host.

**Table 1-14. Memory Performance Enhancement Advice (Continued)**

#	Resolution
4	<p>If the cluster is not a DRS cluster, enable DRS. To enable DRS, perform the following tasks:</p> <ol style="list-style-type: none"> <li>1 Select the cluster, and click the <b>Configure</b> tab.</li> <li>2 Under <b>Services</b>, click <b>vSphere DRS</b>.</li> <li>3 click <b>Edit</b>.</li> </ol> <p>An Edit Cluster Settings dialog box opens.</p> <ol style="list-style-type: none"> <li>4 Click <b>Turn ON vSphere DRS</b>, and click <b>OK</b>.</li> </ol>
5	<p>If the cluster is a DRS cluster:</p> <ul style="list-style-type: none"> <li>■ Increase the number of hosts, and migrate one or more virtual machines to the new host.</li> <li>■ Check the aggressiveness threshold. If the value is low, increase the threshold. It might help avoid hot spots in the cluster.</li> </ul>
6	<p>Add more physical memory to one or more hosts.</p>

## Memory (MB)

The Memory (MB) chart displays memory data counters for clusters. The chart appears at all collection levels except level 1.

### Description

This chart is located in the **Home** view of the cluster **Performance** tab.

**Note** These data counter definitions are for hosts. At the cluster level, the values are collected and totaled. The counter values in the chart represent the aggregate amounts of the host data. The counters that appear in the chart depend on the collection level set for your vCenter Server.

**Table 1-15. Data Counters**

Chart Label	Description
Active	<p>Sum of the active guest physical memory of all powered on virtual machines on the host, plus memory used by basic VMkernel applications. Active memory is estimated by the VMkernel.</p> <ul style="list-style-type: none"> <li>■ Counter: active</li> <li>■ Stats Type: Absolute</li> <li>■ Unit: Megabytes (MB)</li> <li>■ Rollup Type: Average (Minimum/Maximum)</li> <li>■ Collection Level: 2 (4)</li> </ul>
Balloon	<p>Sum of the guest physical memory reclaimed by the balloon driver for all powered on virtual machines on the host.</p> <ul style="list-style-type: none"> <li>■ Counter: vmmemctl</li> <li>■ Stats Type: Absolute</li> <li>■ Unit: Megabytes (MB)</li> <li>■ Rollup Type: Average (Minimum/Maximum)</li> <li>■ Collection Level: 1 (4)</li> </ul>

**Table 1-15. Data Counters (Continued)**

Chart Label	Description
Consumed	<p>Amount of machine memory used on the host.</p> <p>Consumed memory includes virtual machine memory, service console memory, and VMkernel memory.</p> <p>consumed memory = total host memory - free host memory</p> <ul style="list-style-type: none"> <li>■ Counter: consumed</li> <li>■ Stats Type: Absolute</li> <li>■ Unit: Megabytes (MB)</li> <li>■ Rollup Type: Average (Minimum/Maximum)</li> <li>■ Collection Level: 1 (4)</li> </ul>
Granted	<p>Sum of the guest physical memory granted for all powered on virtual machines. Granted memory is mapped to the host's machine memory.</p> <p>Granted memory for a host includes the shared memory of each virtual machine on the host.</p> <ul style="list-style-type: none"> <li>■ Counter: granted</li> <li>■ Stats Type: Absolute</li> <li>■ Unit: Megabytes (MB)</li> <li>■ Rollup Type: Average (Minimum/Maximum)</li> <li>■ Collection Level: 2 (4)</li> </ul>
Swap Used	<p>Sum of the memory swapped by all powered on virtual machines on the host.</p> <ul style="list-style-type: none"> <li>■ Counter: swapused</li> <li>■ Stats Type: Absolute</li> <li>■ Unit: Megabytes (MB)</li> <li>■ Rollup Type: Average (Minimum/Maximum)</li> <li>■ Collection Level: 2 (4)</li> </ul>
Total	<p>Aggregate total memory available to the cluster.</p> <ul style="list-style-type: none"> <li>■ Counter: totalmb</li> <li>■ Stats Type: Absolute</li> <li>■ Unit: Megabytes (MB)</li> <li>■ Rollup Type: Average (Minimum/Maximum)</li> <li>■ Collection Level: 1 (4)</li> </ul>

### Chart Analysis

To ensure best performance, the host memory must be large enough to accommodate the active memory of the virtual machines. The active memory can be smaller than the virtual machine memory size. It allows you to over-provision memory, but still ensures that the virtual machine active memory is smaller than the host memory.

Transient high-usage values usually do not cause performance degradation. For example, memory usage can be high when several virtual machines are started at the same time or when a spike occurs in virtual machine workload. However, a consistently high memory usage value (94% or greater) indicates that the host is probably lacking the memory required to meet the demand. If the active memory size is the same as the granted memory size, the demand for memory is greater than the memory resources available. If the active memory is consistently low, the memory size might be too large.

If the memory usage value is high, and the host has high ballooning or swapping, check the amount of free physical memory on the host. A free memory value of 6% or less indicates that the host cannot handle the demand for memory. It leads to memory reclamation, which might degrade performance.

If the host has enough free memory, check the resource shares, reservation, and limit settings of the virtual machines and resource pools on the host. Verify that the host settings are adequate and not lower than those set for the virtual machines.

If the host has little free memory available, or if you notice a degradation in performance, consider taking the following actions.

**Table 1-16. Memory Performance Enhancement Advice**

#	Resolution
1	Verify that VMware Tools is installed on each virtual machine. The balloon driver is installed with VMware Tools and is critical to performance.
2	Verify that the balloon driver is enabled. The VMkernel regularly reclaims unused virtual machine memory by ballooning and swapping. Generally, it does not impact virtual machine performance.
3	Reduce the memory space on the virtual machine, and correct the cache size if it is too large. This frees up memory for other virtual machines.
4	If the memory reservation of the virtual machine is set to a value much higher than its active memory, decrease the reservation setting so that the VMkernel can reclaim the idle memory for other virtual machines on the host.
5	Migrate one or more virtual machines to a host in a DRS cluster.
6	Add physical memory to the host.

## Memory Consumed

The Memory Consumed chart displays memory usage for the 10 child objects in the cluster with the most consumed memory.

For resource pools and virtual machines in a cluster, this chart is located in the **Resource Pools & Virtual Machines** view of the cluster **Performance** tab. For hosts in a cluster, this chart is located in the **Hosts** view of the cluster **Performance** tab.

**Table 1-17. Data Counters**

Chart Label	Description
<i>resource_pool,</i> <i>virtual_machine,</i> or <i>host</i>	<p>Amount of machine memory used by all resource pools and virtual machines in the cluster or by all hosts in the cluster, depending on the cluster view.</p> <p>Consumed memory includes virtual machine memory, service console memory, and VMkernel memory.</p> <p>consumed memory = total host memory - free host memory</p> <ul style="list-style-type: none"> <li>■ Counter: consumed</li> <li>■ Stats Type: Absolute</li> <li>■ Unit: MegaBytes (MB)</li> <li>■ Rollup Type: Average (Minimum/Maximum)</li> <li>■ Collection Level: 1 (4)</li> </ul>

## Chart Analysis

Memory usage is not an indicator of performance problems. Memory can be high if a host is swapping or ballooning, which can result in virtual machine guest swapping. In such cases, check for other problems, such as CPU over-commitment or storage latencies.

If you have constantly high memory usage in a cluster, resource pool, or vApp, consider taking the following actions.

**Table 1-18. Memory Performance Enhancement Advice**

#	Resolution
1	Verify that VMware Tools is installed on each virtual machine.
2	Verify that the balloon driver is enabled. The balloon driver is installed with VMware Tools and is critical to performance. The VMkernel regularly reclaims unused virtual machine memory by ballooning and swapping. Generally, it does not impact virtual machine performance.
3	If the balloon value is high, check the resource shares, reservations, and limits for the virtual machines and resource pools on the hosts. Verify that the host's settings are adequate and not lower than those set for the virtual machine. If free memory is available on the hosts and the virtual machines are experiencing high swap or balloon memory, the virtual machine (or resource pool, if it belongs to one) has reached its resource limit. Check the maximum resource limit set on that host.
4	If the cluster is not a DRS cluster, enable DRS. To enable DRS, perform the following tasks: <ol style="list-style-type: none"> <li>1 Select the cluster, and click the <b>Configure</b> tab.</li> <li>2 Under <b>Services</b>, click <b>vSphere DRS</b>.</li> <li>3 click <b>Edit</b>.  An Edit Cluster Settings dialog box opens.</li> <li>4 Click <b>Turn ON vSphere DRS</b>, and click <b>OK</b>.</li> </ol>
5	If the cluster is a DRS cluster: <ul style="list-style-type: none"> <li>■ Increase the number of hosts, and migrate one or more virtual machines to the new host.</li> <li>■ Check the aggressiveness threshold. If the value is low, increase the threshold. It might help avoid hot spots in the cluster.</li> </ul>
6	Add more physical memory to one or more hosts.

## Network (Mbps)

The Network (Mbps) chart displays network speed for the 10 hosts in the cluster with the most network usage.

This chart is located in the **Hosts** view of the Cluster **Performance** tab.

**Table 1-19. Data Counters**

Chart Label	Description
<host>	Average rate at which data is transmitted and received across all NIC instances on the host. <ul style="list-style-type: none"> <li>■ Counter: usage</li> <li>■ Stats Type: Rate</li> <li>■ Unit: Megabits per second (Mbps)</li> <li>■ Rollup Type: Average (Minimum/Maximum)</li> <li>■ Collection Level: 1 (4)</li> </ul>

## Chart Analysis

Network performance depends on the application workload and network configuration. Dropped network packets indicate a bottleneck in the network. To determine whether packets are being dropped, use `esxtop` or the advanced performance charts to examine the `droppedTx` and `droppedRx` network counter values.

If packets are being dropped, adjust the virtual machine shares. If packets are not being dropped, check the size of the network packets and the data receive and transfer rates. In general, the larger the network packets, the faster the network speed. When the packet size is large, fewer packets are transferred, which reduces the amount of CPU required to process the data. When network packets are small, more packets are transferred but the network speed is slower because more CPU is required to process the data.

---

**Note** In some instances, large packets might result in a high network latency. To check the network latency, use the VMware AppSpeed performance monitoring application or a third-party application.

---

If packets are not being dropped and the data receive rate is slow, the host is probably lacking the CPU resources required to handle the load. Check the number of virtual machines assigned to each physical NIC. If necessary, perform load balancing by moving virtual machines to different vSwitches or by adding more NICs to the host. You can also move virtual machines to another host or increase the host CPU or virtual machine CPU.

If you experience network-related performance problems, also consider taking the following actions.

**Table 1-20. Networking Performance Enhancement Advice**

#	Resolution
1	Verify that VMware Tools is installed on each virtual machine.
2	If possible, use <code>vmxnet3</code> NIC drivers, which are available with VMware Tools. They are optimized for high performance.
3	If virtual machines running on the same host communicate with each other, connect them to the same vSwitch to avoid transferring packets over the physical network.
4	Assign each physical NIC to a port group and a vSwitch.
5	Use separate physical NICs to handle the different traffic streams, such as network packets generated by virtual machines, iSCSI protocols, vMotion tasks.
6	Ensure that the physical NIC capacity is large enough to handle the network traffic on that vSwitch. If the capacity is not enough, consider using a high-bandwidth physical NIC (10 Gbps). Alternatively, consider moving some virtual machines to a vSwitch with a lighter load or to a new vSwitch.
7	If packets are being dropped at the vSwitch port, increase the virtual network driver ring buffers where applicable.
8	Verify that the reported speed and duplex settings for the physical NIC match the hardware expectations and that the hardware is configured to run at its maximum capability. For example, verify that NICs with 1 Gbps are not reset to 100 Mbps because they are connected to an older switch.
9	Verify that all NICs are running in full duplex mode. Hardware connectivity problems might result in a NIC resetting itself to a lower speed or half duplex mode.
10	Use vNICs that are TCP Segmentation Offload (TSO)-capable, and verify that TSO-Jumbo Frames are enabled where possible.

## Data centers

The data center charts contain information about CPU, disk, memory, and storage usage for data centers. The help topic for each chart contains information about the data counters displayed in that chart. The counters available are determined by the collection level set for vCenter Server.

### CPU (MHz)

The CPU (MHz) chart displays CPU usage for the 10 clusters in the data center with the most CPU usage.

This chart is located in the Clusters view of the Datacenters **Performance** tab.

**Table 1-21. Data Counters**

Chart Label	Description
<cluster>	<p>Amount of CPU currently in use by the cluster. The active CPU usage is approximately equal to the ratio of the used CPU cycles to the available CPU cycles.</p> <p>The maximum possible value is the frequency of the processors multiplied by the number of cores. For example, a two-way SMP virtual machine using 4000MHz on a host that has four 2GHz processors is using 50% of the CPU (<math>4000 \div 4 \times 2000 = 0.5</math>).</p> <ul style="list-style-type: none"> <li>■ Counter: usagemhz</li> <li>■ Stats Type: Rate</li> <li>■ Unit: MegaHertz (MHz)</li> <li>■ Rollup Type: Average (Minimum/Maximum)</li> <li>■ Collection Level: 1 (4)</li> </ul>

### Chart Analysis

A short spike in CPU usage indicates that you are making the best use of cluster resources. However, if the value is constantly high, the CPU demanded is likely greater than the CPU capacity available. A high CPU usage value can lead to increased ready time and processor queuing of the virtual machines on the hosts in the cluster.

If performance is impacted, consider taking the following actions.

**Table 1-22. CPU Performance Enhancement Advice**

#	Resolution
1	Verify that VMware Tools is installed on each virtual machine.
2	If the cluster is not a DRS cluster, enable DRS. To enable DRS, perform the following tasks: <ol style="list-style-type: none"> <li>1 Select the cluster, and click the <b>Configure</b> tab.</li> <li>2 Under <b>Services</b>, click <b>vSphere DRS</b>.</li> <li>3 click <b>Edit</b>.</li> </ol> <p>An Edit Cluster Settings dialog box opens.</p> <ol style="list-style-type: none"> <li>4 Click <b>Turn ON vSphere DRS</b>, and click <b>OK</b>.</li> </ol>



**Table 1-22. CPU Performance Enhancement Advice (Continued)**

#	Resolution
3	If the cluster is a DRS cluster: <ul style="list-style-type: none"> <li>■ Increase the number of hosts, and migrate one or more virtual machines to the new host.</li> <li>■ Check the aggressiveness threshold. If the value is low, increase the threshold. This might help avoid hot spots in the cluster.</li> </ul>
4	Migrate one or more virtual machines to a new host.
5	Upgrade the physical CPUs or cores on each host in the cluster if necessary.
6	Enable CPU-saving features, such as TCP Segmentation Offload.
7	Replace software I/O with the dedicated hardware, such as iSCSI HBAs or TCP Segmentation Offload NICs.

## Memory (MB)

The Memory (MB) chart displays the average amount of consumed memory for the 10 clusters in the data center with the most consumed memory.

This chart is located in the **Clusters** view of the Datacenters **Performance** tab.

**Table 1-23. Data Counters**

Chart Label	Description
<cluster>	Amount of host machine memory used by all powered on virtual machines in the cluster. <ul style="list-style-type: none"> <li>■ Counter: consumed</li> <li>■ Stats Type: Absolute</li> <li>■ Unit: MegaBytes (MB)</li> <li>■ Rollup Type: Average (Minimum/Maximum)</li> <li>■ Collection Level: 1 (4)</li> </ul>

## Chart Analysis

A cluster's consumed memory consists of virtual machine consumed memory and overhead memory. It does not include host-specific overhead memory, such as memory used by the service console or VMkernel.

If you experience problems with cluster memory usage, use the thumbnail cluster charts to examine memory usage for each cluster and increase memory resources if needed.

If the cluster is a DRS cluster, check the aggressiveness threshold. If the value is low, increase the threshold. Increasing the threshold might help avoid hot spots in the cluster.

## Space used by Datastore in GB

The Space used by Datastore in GB chart displays the 10 datastores in the data center with the most used disk space.

This chart is located in the **Storage** view of the Datacenter **Performance** tab.

**Table 1-24. Data Counters**

Chart Label	Description
<datastore>	<p>Amount of used storage space on the 10 datastores with the most used space.</p> <ul style="list-style-type: none"> <li>■ Counter: used</li> <li>■ Stats Type: Absolute</li> <li>■ Unit: GigaBytes (GB)</li> <li>■ Rollup Type: Latest</li> <li>■ Collection Level: 1</li> </ul>

### Chart Analysis

The datastore is at full capacity when the used space is equal to the capacity. Allocated space can be larger than datastore capacity, for example, when you have snapshots and thin-provisioned disks. If possible, you can provision more space to the datastore, or you can add disks to the datastore or use shared datastores.

If snapshot files are consuming high datastore space, consider consolidating them to the virtual disk when they are no longer needed. Consolidating the snapshots deletes the redo log files and removes the snapshots from the vSphere Web Client user interface. For information about consolidating the data center, see the vSphere documentation.

### Space Utilization By File Type

The Space Utilization By File Type chart displays datastore space usage for virtual disks, swap files, snapshot files, and other virtual machine files.

---

**Note** This chart does not show historical statistics. It only shows the most recently available data, which may be up to 30 minutes late, depending on when the last statistics rollup occurred. In addition, statistics are not collected across all datastores at one time. They are collected asynchronously.

---

The Space Utilization by File Type chart is located in the **Storage** view of the data center **Performance** tab.

## Datastore Counters

**Table 1-25. Data Counters**

File Type	Description
Virtual Disks	<p>Amount of disk space used by virtual disk files.</p> <p>Virtual disk files store the contents of the virtual machine's hard disk drive. It includes information that you write to a virtual machine's hard disk, such as the operating system, program files, and data files. The files have the extension <code>.vmdk</code> and appear as a physical disk drive to a guest operating system.</p> <hr/> <p><b>Note</b> Delta disks, which also have an extension <code>.vmdk</code>, are not included in this file type.</p> <hr/> <ul style="list-style-type: none"> <li>■ Counter: used</li> <li>■ Stats Type: Absolute</li> <li>■ Unit: Gigabytes (GB)</li> <li>■ Rollup Type: Latest</li> <li>■ Collection Level: 1 (4)</li> </ul>
Swap Files	<p>Amount of disk space used by swap files.</p> <p>Swap files back up the virtual machine's physical memory.</p> <ul style="list-style-type: none"> <li>■ Counter: used</li> <li>■ Stats Type: Absolute</li> <li>■ Unit: Gigabytes (GB)</li> <li>■ Rollup Type: Latest</li> <li>■ Collection Level: 1 (4)</li> </ul>
Snapshots	<p>Amount of disk space used by virtual machine snapshot files.</p> <p>Snapshot files store information about virtual machine snapshots. They include snapshot state files and delta disk files. A snapshot state file stores the running state of the virtual machine at the time of the snapshot. It has the extension <code>.vmsn</code>. A delta disk file stores the updates made by the virtual machine to the virtual disks after a snapshot is taken.</p> <ul style="list-style-type: none"> <li>■ Counter: used</li> <li>■ Stats Type: Absolute</li> <li>■ Unit: Gigabytes (GB)</li> <li>■ Rollup Type: Latest</li> <li>■ Collection Level: 1 (4)</li> </ul>
Other VM Files	<p>Amount of disk space used by all other virtual machine files, such as configuration files and log files.</p> <ul style="list-style-type: none"> <li>■ Counter: used</li> <li>■ Stats Type: Absolute</li> <li>■ Unit: Gigabytes (GB)</li> <li>■ Rollup Type: Latest</li> <li>■ Collection Level: 1 (4)</li> </ul>
Other	<p>Amount of disk space used by all other non-virtual machine files, such as documentation files and backup files.</p>
Free Space	<p>Amount of disk space not currently in use.</p>
Total Space	<p>Amount of disk space available to the datastore. It defines the datastore capacity. The chart displays the information for datastores but not for data centers.</p> <p>total space = virtual disk space + swap file space + snapshot space + other VM file space + other space + free space</p>

## Chart Analysis

The datastore is at full capacity when the used space is equal to the capacity. Allocated space can be larger than datastore capacity, for example, when you have snapshots and thin-provisioned disks. If possible, you can provision more space to the datastore, or you can add disks to the datastore or use shared datastores.

If snapshot files are consuming high datastore space, consider consolidating them to the virtual disk when they are no longer needed. Consolidating the snapshots deletes the redo log files and removes the snapshots from the vSphere Web Client user interface. For information about consolidating the data center, see the vSphere documentation.

## Datastores and Datastore Clusters

The datastore charts contain information about disk usage for datastores or the datastores that are part of a cluster. The help topic for each chart contains information about the data counters displayed in that chart. The counters available are determined by the collection level set for vCenter Server.

### Space in GB

The Space in GB chart displays space usage data counters for datastores.

This chart is located in the **Space** view of the datastore or datastore cluster **Performance** tab.

**Table 1-26. Data Counters**

Chart Label	Description
Allocated	<p>Amount of physical space provisioned by an administrator for the datastore. It is the storage size up to which files on the datastore can grow. Allocated space is not always in use.</p> <ul style="list-style-type: none"> <li>■ Counter: provisioned</li> <li>■ Stats Type: Absolute</li> <li>■ Unit: Gigabytes (GB)</li> <li>■ Rollup Type: Latest</li> <li>■ Collection Level: 1</li> </ul>
Used	<p>Amount of physical datastore space in use.</p> <ul style="list-style-type: none"> <li>■ Counter: used</li> <li>■ Stats Type: Absolute</li> <li>■ Unit: Gigabytes (GB)</li> <li>■ Rollup Type: Latest</li> <li>■ Collection Level: 1</li> </ul>
Capacity	<p>Maximum capacity of the datastore.</p> <p>capacity = virtual machine file space + non-virtual machine file space + free space</p> <p><b>Note</b> Storage data is collected and updated in the overview charts every 30 minutes. Therefore, if you refresh the datastore, the capacity value might only be updated in the datastore <b>Summary</b> tab, and not in the overview charts.</p> <ul style="list-style-type: none"> <li>■ Counter: capacity</li> <li>■ Stats Type: Absolute</li> <li>■ Unit: Gigabytes (GB)</li> <li>■ Rollup Type: Latest</li> <li>■ Collection Level: 1</li> </ul>

### Chart Analysis

The datastore is at full capacity when the used space is equal to the capacity. Allocated space can be larger than datastore capacity, for example, when you have snapshots and thin-provisioned disks. If possible, you can provision more space to the datastore, or you can add disks to the datastore or use shared datastores.

If snapshot files are consuming high datastore space, consider consolidating them to the virtual disk when they are no longer needed. Consolidating the snapshots deletes the redo log files and removes the snapshots from the vSphere Web Client user interface. For information about consolidating the data center, see the vSphere documentation.

### Space Utilization By File Type

The Space Utilization by File Type chart displays space used by virtual disks, swap files, snapshot files, and other virtual machine files on the datastore or the datastore cluster.

**Note** This chart does not show historical statistics. It only shows the most recently available data, which may be up to 30 minutes late, depending on when the last statistics rollup occurred. In addition, statistics are not collected across all datastores at one time. They are collected asynchronously.

The Space Utilization by File Type chart is located in the **Storage** view of the datastore **Performance** tab. The counters can also be displayed for datastore cluster charts.

## Datastore Counters

**Table 1-27. Data Counters**

File Type	Description
Virtual Disks	<p>Amount of disk space used by virtual disk files.</p> <p>Virtual disk files store the contents of the virtual machine's hard disk drive. It includes information that you write to a virtual machine's hard disk, such as the operating system, program files, and data files. The files have the extension <code>.vmdk</code> and appear as a physical disk drive to a guest operating system.</p> <hr/> <p><b>Note</b> Delta disks, which also have an extension <code>.vmdk</code>, are not included in this file type.</p> <hr/> <ul style="list-style-type: none"> <li>■ Counter: used</li> <li>■ Stats Type: Absolute</li> <li>■ Unit: Gigabytes (GB)</li> <li>■ Rollup Type: Latest</li> <li>■ Collection Level: 1 (4)</li> </ul>
Swap Files	<p>Amount of disk space used by swap files.</p> <p>Swap files back up the virtual machine's physical memory.</p> <ul style="list-style-type: none"> <li>■ Counter: used</li> <li>■ Stats Type: Absolute</li> <li>■ Unit: Gigabytes (GB)</li> <li>■ Rollup Type: Latest</li> <li>■ Collection Level: 1 (4)</li> </ul>
Snapshots	<p>Amount of disk space used by virtual machine snapshot files.</p> <p>Snapshot files store information about virtual machine snapshots. They include snapshot state files and delta disk files. A snapshot state file stores the running state of the virtual machine at the time of the snapshot. It has the extension <code>.vmsn</code>. A delta disk file stores the updates made by the virtual machine to the virtual disks after a snapshot is taken.</p> <ul style="list-style-type: none"> <li>■ Counter: used</li> <li>■ Stats Type: Absolute</li> <li>■ Unit: Gigabytes (GB)</li> <li>■ Rollup Type: Latest</li> <li>■ Collection Level: 1 (4)</li> </ul>
Other VM Files	<p>Amount of disk space used by all other virtual machine files, such as configuration files and log files.</p> <ul style="list-style-type: none"> <li>■ Counter: used</li> <li>■ Stats Type: Absolute</li> <li>■ Unit: Gigabytes (GB)</li> <li>■ Rollup Type: Latest</li> <li>■ Collection Level: 1 (4)</li> </ul>
Other	<p>Amount of disk space used by all other non-virtual machine files, such as documentation files and backup files.</p>

**Table 1-27. Data Counters (Continued)**

File Type	Description
Free Space	Amount of disk space not currently in use.
Total Space	Amount of disk space available to the datastore. It defines the datastore capacity. The chart displays the information for datastores but not for data centers.  total space = virtual disk space + swap file space + snapshot space + other VM file space + other space + free space

### Chart Analysis

The datastore is at full capacity when the used space is equal to the capacity. Allocated space can be larger than datastore capacity, for example, when you have snapshots and thin-provisioned disks. If possible, you can provision more space to the datastore, or you can add disks to the datastore or use shared datastores.

If snapshot files are consuming high datastore space, consider consolidating them to the virtual disk when they are no longer needed. Consolidating the snapshots deletes the redo log files and removes the snapshots from the vSphere Web Client user interface. For information about consolidating the data center, see the vSphere documentation.

### Space used by Datastore in GB

The Space used by Datastore in GB chart displays the 10 datastores in the data center with the most used disk space.

This chart is located in the **Storage** view of the Datacenter **Performance** tab.

**Table 1-28. Data Counters**

Chart Label	Description
<datastore>	Amount of used storage space on the 10 datastores with the most used space. <ul style="list-style-type: none"> <li>■ Counter: used</li> <li>■ Stats Type: Absolute</li> <li>■ Unit: GigaBytes (GB)</li> <li>■ Rollup Type: Latest</li> <li>■ Collection Level: 1</li> </ul>

### Chart Analysis

The datastore is at full capacity when the used space is equal to the capacity. Allocated space can be larger than datastore capacity, for example, when you have snapshots and thin-provisioned disks. If possible, you can provision more space to the datastore, or you can add disks to the datastore or use shared datastores.

If snapshot files are consuming high datastore space, consider consolidating them to the virtual disk when they are no longer needed. Consolidating the snapshots deletes the redo log files and removes the snapshots from the vSphere Web Client user interface. For information about consolidating the data center, see the vSphere documentation.

## Space Utilization by Virtual Machine

The Space Utilization by Virtual Machine chart displays the amount of space used by the five virtual machines with the most space used on the datastore or the datastores in the cluster.

**Note** This chart does not show historical statistics. It only shows the most recently available data, which may be up to 30 minutes late, depending on when the last statistics rollup occurred. In addition, statistics are not collected across all datastores at one time. They are collected asynchronously.

The Space Utilization by Virtual Machine chart is located in the **Space** view of the datastore **Performance** tab. The counter can also be displayed for datastore cluster charts.

**Table 1-29. Data Counters**

Chart Label	Description
<i>virtual_machine</i>	<p>Amount of datastore space used by the five virtual machines with the most used datastore space.</p> <ul style="list-style-type: none"> <li>■ Counter: used</li> <li>■ Stats Type: Absolute</li> <li>■ Unit: Gigabytes (GB)</li> <li>■ Rollup Type: Latest</li> <li>■ Collection Level: 1</li> </ul>

### Chart Analysis

The datastore is at full capacity when the used space is equal to the capacity. Allocated space can be larger than datastore capacity, for example, when you have snapshots and thin-provisioned disks. If possible, you can provision more space to the datastore, or you can add disks to the datastore or use shared datastores.

If snapshot files are consuming high datastore space, consider consolidating them to the virtual disk when they are no longer needed. Consolidating the snapshots deletes the redo log files and removes the snapshots from the vSphere Web Client user interface. For information about consolidating the data center, see the vSphere documentation.

## Space Allocated by Datastore in GB

The Space allocated by Datastore in GB displays the top 10 datastores, virtual machines in the datastore-cluster with most provisioned space.

This chart is located in the **Space** view of the Datacenter **Performance** tab.



**Table 1-30. Data Counters**

Chart Label	Description
<datastore>	<p>Amount of provisioned storage space on the top 10 datastores with the most provisioned space.</p> <ul style="list-style-type: none"> <li>■ Counter: provisioned</li> <li>■ Stats Type: Absolute</li> <li>■ Unit: KiloBytes (KB)</li> <li>■ Rollup Type: Latest</li> <li>■ Collection Level: 1</li> </ul>

### Chart Analysis

The datastore is at full capacity when the used space is equal to the capacity. Allocated space can be larger than datastore capacity, for example, when you have snapshots and thin-provisioned disks. If possible, you can provision more space to the datastore, or you can add disks to the datastore or use shared datastores.

If snapshot files are consuming high datastore space, consider consolidating them to the virtual disk when they are no longer needed. Consolidating the snapshots deletes the redo log files and removes the snapshots from the vSphere Web Client user interface. For information about consolidating the data center, see the vSphere documentation.

### Space Capacity by Datastore in GB

The Space capacity by Datastore in GB displays the top 10 configured size of the datastores in the datastore cluster.

This chart is located in the **Space** view of the Datacenter **Performance** tab.

**Table 1-31. Data Counters**

Chart Label	Description
<datastore>	<p>Configured size of the datastores in the datastore cluster.</p> <ul style="list-style-type: none"> <li>■ Counter: capacity</li> <li>■ Stats Type: Absolute</li> <li>■ Unit: KiloBytes (KB)</li> <li>■ Rollup Type: Latest</li> <li>■ Collection Level: 1</li> </ul>

### Chart Analysis

The datastore is at full capacity when the used space is equal to the capacity. Allocated space can be larger than datastore capacity, for example, when you have snapshots and thin-provisioned disks. If possible, you can provision more space to the datastore, or you can add disks to the datastore or use shared datastores.

If snapshot files are consuming high datastore space, consider consolidating them to the virtual disk when they are no longer needed. Consolidating the snapshots deletes the redo log files and removes the snapshots from the vSphere Web Client user interface. For information about consolidating the data center, see the vSphere documentation.

## Storage I/O Control Normalized Latency

This chart displays the normalized latency in microseconds on the datastore. Storage I/O Control monitors latency to detect congestion on the datastore. This metric computes a weighted response time across all hosts and VMs accessing the datastore. I/O count is used as the weight for the response time. It captures the device level latency and does not include any queuing inside the hypervisor storage stack or inside the VM. It is adjusted for the I/O size. High latencies that are the result of large I/Os are discounted so as not to make the datastore seem slower than it really is. Data for all virtual machines is combined. This chart displays zero values when Storage I/O Control is disabled.

This chart is located in the **Performance** view of the datastore **Performance** tab. The `sizeNormalizedDatastoreLatency` counter can also be displayed for datastore cluster charts.

**Table 1-32. Data Counters**

Chart Label	Description
Storage I/O Control Normalized Latency	<p>Storage I/O Control monitors latency to detect congestion on the datastore.</p> <ul style="list-style-type: none"> <li>■ Counter: <code>sizeNormalizedDatastoreLatency</code></li> <li>■ Stats Type: Absolute</li> <li>■ Unit: Microseconds</li> <li>■ Rollup Type: Average</li> <li>■ Collection Level: 3</li> </ul>

## Storage I/O Control Aggregate IOPs

This chart displays the number of I/O operations per one second on the datastore, aggregated across all hosts, and virtual machines accessing this datastore. The chart displays zero values when Storage I/O Control is disabled.

This chart is located in the **Performance** view of the datastore or the datastore cluster **Performance** tab. The counter can be displayed for datastore and datastore cluster charts.

**Table 1-33. Data Counters**

Chart Label	Description
Storage I/O Control Aggregate IOPs	<p>Number of I/O operations per second on the datastore, aggregated across all hosts, and virtual machines accessing the datastore.</p> <ul style="list-style-type: none"> <li>■ Counter: <code>datastoreIops</code></li> <li>■ Stats Type: Absolute</li> <li>■ Unit: Number</li> <li>■ Rollup Type: Average</li> <li>■ Collection Level: 3</li> </ul>

## Storage I/O Control Activity

This chart displays the percentage of time during which Storage I/O Control actively controlled latency of the datastore.

This chart is located in the **Performance** views of the datastore **Performance** tabs. The counter can also be displayed for datastore cluster charts.

**Table 1-34. Data Counters**

Chart Label	Description
Storage I/O Control Activity	<p>This is the percentage of time during which the Storage I/O Control actively controlled the I/O latency for the datastore.</p> <ul style="list-style-type: none"> <li>■ Counter: siocActiveTimePercentage</li> <li>■ Stats Type: Absolute</li> <li>■ Unit: Percent</li> <li>■ Rollup Type: Average</li> <li>■ Collection Level: 3</li> </ul>

## Average Device Latency per Host

This chart displays the average amount of latency on a host device. The chart displays the 10 hosts with the highest device latency.

This chart is located in the **Performance** view of the datastore **Performance** tab.

**Table 1-35. Data Counters**

Chart Label	Description
Average Device Latency per Host	<p>Measures the amount of time, in milliseconds, to complete a SCSI command issued from the physical device.</p> <ul style="list-style-type: none"> <li>■ Counter: deviceLatency</li> <li>■ Stats Type: Absolute</li> <li>■ Unit: Milliseconds (ms)</li> <li>■ Rollup Type: Average</li> <li>■ Collection Level: 3</li> </ul>

## Maximum Queue Depth per Host

This chart displays maximum queue depth that hosts are currently maintaining for the datastore. When Storage I/O is enabled, queue depth can change over time when congestion is detected at the array.

This chart is located in the **Performance** view of the datastore **Performance** tab. The chart displays information about the ten hosts with the highest values.

**Table 1-36. Data Counters**

Chart Label	Description
Max Queue Depth per Host	<p>Maximum queue depth. Queue depth is the number of commands the SCSI driver queues to the HBA.</p> <ul style="list-style-type: none"> <li>■ Counter: maxQueueDepth</li> <li>■ Stats Type: Absolute</li> <li>■ Unit: Number</li> <li>■ Rollup Type: Average</li> <li>■ Collection Level: 3</li> </ul>

## Read IOPs per Host

This chart displays the per-host disk read rates for a datastore. The chart displays information about the ten hosts with the highest values.

This chart is located in the **Performance** view of the datastore **Performance** tab.

**Table 1-37. Data Counters**

Chart Label	Description
Read IOPs per Host	<p>Number of disk read commands completed on each disk on the host, per second.</p> <p>Read rate = blocks read per second × block size</p> <ul style="list-style-type: none"> <li>■ Counter: numberReadAveraged</li> <li>■ Stats Type: Rate</li> <li>■ Unit: Number</li> <li>■ Rollup Type: Average</li> <li>■ Collection Level: 3</li> </ul>

## Write IOPs Per Host

This chart displays the per-host disk write rates for a datastore. The chart displays information about the 10 hosts with the highest values.

This chart is located in the **Performance** view of the datastore **Performance** tab.

**Table 1-38. Data Counters**

Chart Label	Description
Write IOPs per Host	<p>Number of disk write commands completed on each disk on the host, per second.</p> <p>Write rate = blocks written per second × block size</p> <ul style="list-style-type: none"> <li>■ Counter: numberWriteAveraged</li> <li>■ Stats Type: Rate</li> <li>■ Unit: Number</li> <li>■ Rollup Type: Average</li> <li>■ Collection Level: 3</li> </ul>

## Average Read Latency per Virtual Machine Disk

This chart displays the top ten virtual machine disks with the highest average read latency in milliseconds. Data is not displayed when the virtual machine is powered off.

This chart is located in the **Performance** view of the datastore **Performance** tab. The counter can also be displayed for datastore cluster charts.

**Table 1-39. Data Counters**

Chart Label	Description
Average Read Latency per Virtual Machine Disk	<p>Latency measures the time used to process a SCSI command issued by the guest OS to the virtual machine. The kernel latency is the time VMkernel takes to process an I/O request. The device latency is the time it takes the hardware to handle the request.</p> <p>Total latency = kernelLatency + deviceLatency.</p> <ul style="list-style-type: none"> <li>■ Counter: totalReadLatency</li> <li>■ Stats Type: Absolute</li> <li>■ Unit: Milliseconds (ms)</li> <li>■ Rollup Type: Average</li> <li>■ Collection Level: 3</li> </ul>

## Average Write Latency per Virtual Machine Disk

This chart displays the top ten virtual machine disks with the highest average write latency in milliseconds. Data is not displayed when the virtual machine is powered off.

This chart is located in the **Performance** view of the datastore **Performance** tab. The counter can also be displayed for datastore cluster charts.

**Table 1-40. Data Counters**

Chart Label	Description
Average Write Latency per Virtual Machine Disk	<p>Latency measures the time used to process a SCSI command issued by the guest OS to the virtual machine. The kernel latency is the time VMkernel takes to process an I/O request. The device latency is the time it takes the hardware to handle the request.</p> <p>Total latency = kernelLatency + deviceLatency.</p> <ul style="list-style-type: none"> <li>■ Counter: totalWriteLatency</li> <li>■ Stats Type: Absolute</li> <li>■ Unit: Milliseconds (ms)</li> <li>■ Rollup Type: Average</li> <li>■ Collection Level: 3</li> </ul>

## Read IOPs per Virtual Machine Disk

This chart displays the top ten virtual machines with the highest number of read operations. Data is not displayed when the virtual machine is powered off.

This chart is located in the **Performance** view of the datastore **Performance** tab. The counter can also be displayed for datastore cluster charts.

**Table 1-41. Data Counters**

Chart Label	Description
Read IOPs per Virtual Machine Disk	<p>Number of disk read commands completed on each virtual machine disk, per second.</p> <p>Read rate = blocks read per second × block size</p> <ul style="list-style-type: none"> <li>■ Counter: numberReadAveraged</li> <li>■ Stats Type: Rate</li> <li>■ Unit: Number</li> <li>■ Rollup Type: Average</li> <li>■ Collection Level: 3</li> </ul>

## Write IOPs Per Virtual Machine Disk

This chart displays the 10 virtual machines with the highest number of write operations. Data is not displayed when the virtual machine is powered off.

This chart is located in the **Performance** view of the datastore **Performance** tab. The counter can also be displayed for datastore cluster charts.

**Table 1-42. Data Counters**

Chart Label	Description
Write IOPs per Virtual Machine Disk	<p>Number of disk write commands completed on each virtual machine disk on the host.</p> <p>Write rate = blocks read per second × block size</p> <ul style="list-style-type: none"> <li>■ Counter: numberWriteAveraged</li> <li>■ Stats Type: Rate</li> <li>■ Unit: Number</li> <li>■ Rollup Type: Average</li> <li>■ Collection Level: 3</li> </ul>

## Virtual Machine Observed Latency per Datastore

This chart displays the average datastore latency as observed by the virtual machines.

This chart is located in the **Performance** view of the datastore cluster **Performance** tab.

**Table 1-43. Data Counters**

Chart Label	Description
VM observed latency report per Datastore	<p>This is the average datastore latency as observed by the virtual machines in the datastore cluster.</p> <ul style="list-style-type: none"> <li>■ Counter: datastoreVMObservedLatency</li> <li>■ Stats Type: Absolute</li> <li>■ Unit: Microseconds</li> <li>■ Rollup Type: Latest</li> <li>■ Collection Level: 3</li> </ul>

## Hosts

The hosts charts contain information about CPU, disk, memory, network, and storage usage for hosts. The help topic for each chart contains information about the data counters displayed in that chart. The counters available are determined by the collection level set for vCenter Server.

### CPU (%)

The CPU (%) chart displays CPU usage for the host.

This chart is located in the Home view of the Host **Performance** tab.

**Table 1-44. Data Counters**

Chart Label	Description
Usage	<p>Actively used CPU, as a percentage of the total available CPU, for each physical CPU on the host.</p> <p>Active CPU is approximately equal to the ratio of the used CPU to the available CPU.</p> <p>Available CPU = # of physical CPUs × clock rate.</p> <p>100% represents all CPUs on the host. For example, if a four-CPU host is running a virtual machine with two CPUs, and the usage is 50%, the host is using two CPUs completely.</p> <ul style="list-style-type: none"> <li>■ Counter: usage</li> <li>■ Stats Type: Rate</li> <li>■ Unit: Percentage (%)</li> <li>■ Rollup Type: Average (Minimum/Maximum)</li> <li>■ Collection Level: 1 (4)</li> </ul>

### Chart Analysis

A short spike in CPU usage indicates that you are making the best use of the host resources. However, if the value is constantly high, the host is probably lacking the CPU required to meet the demand. A high CPU usage value can lead to increased ready time and processor queuing of the virtual machines on the host.

If performance is impacted, consider taking the following actions.

**Table 1-45. CPU Performance Enhancement Advice**

#	Resolution
1	Verify that VMware Tools is installed on every virtual machine on the host.
2	Set the CPU reservations for all high-priority virtual machines to guarantee that they receive the CPU cycles required.
3	Reduce the number of virtual CPUs on a virtual machine to only the number required to execute the workload. For example, a single-threaded application on a four-way virtual machine only benefits from a single vCPU. But the hypervisor's maintenance of the three idle vCPUs takes CPU cycles that could be used for other work.
4	If the host is not already in a DRS cluster, add it to one. If the host is in a DRS cluster, increase the number of hosts and migrate one or more virtual machines onto the new host.

**Table 1-45. CPU Performance Enhancement Advice (Continued)**

#	Resolution
5	Upgrade the physical CPUs or cores on the host if necessary.
6	Use the newest version of hypervisor software, and enable CPU-saving features such as TCP Segmentation Offload, large memory pages, and jumbo frames.

## CPU (MHz)

The CPU (MHz) chart displays CPU usage for the host.

This chart is located in the Home view of the Host **Performance** tab.

**Table 1-46. Data Counters**

Chart Label	Description
Usage	<p>The sum, in megahertz, of the actively used CPU of all powered on virtual machines on a host.</p> <p>The maximum possible value is the frequency of the processors multiplied by the number of processors. For example, if you have a host with four 2GHz CPUs running a virtual machine that is using 4000MHz, the host is using two CPUs completely.</p> $4000 \div (4 \times 2000) = 0.50$ <ul style="list-style-type: none"> <li>■ Counter: usagemhz</li> <li>■ Stats Type: Rate</li> <li>■ Unit: MegaHertz (MHz)</li> <li>■ Rollup Type: Average (Minimum/Maximum)</li> <li>■ Collection Level: 1 (4)</li> </ul>

## Chart Analysis

A short spike in CPU usage indicates that you are making the best use of the host resources. However, if the value is constantly high, the host is probably lacking the CPU required to meet the demand. A high CPU usage value can lead to increased ready time and processor queuing of the virtual machines on the host.

If performance is impacted, consider taking the following actions.

**Table 1-47. CPU Performance Enhancement Advice**

#	Resolution
1	Verify that VMware Tools is installed on every virtual machine on the host.
2	Set the CPU reservations for all high-priority virtual machines to guarantee that they receive the CPU cycles required.
3	Reduce the number of virtual CPUs on a virtual machine to only the number required to execute the workload. For example, a single-threaded application on a four-way virtual machine only benefits from a single vCPU. But the hypervisor's maintenance of the three idle vCPUs takes CPU cycles that could be used for other work.
4	If the host is not already in a DRS cluster, add it to one. If the host is in a DRS cluster, increase the number of hosts and migrate one or more virtual machines onto the new host.



**Table 1-47. CPU Performance Enhancement Advice (Continued)**

#	Resolution
5	Upgrade the physical CPUs or cores on the host if necessary.
6	Use the newest version of hypervisor software, and enable CPU-saving features such as TCP Segmentation Offload, large memory pages, and jumbo frames.

## CPU Usage

The CPU Usage chart displays CPU usage of the 10 virtual machines on the host with the most CPU usage.

This chart is located in the Virtual Machines view of the host **Performance** tab.

**Table 1-48. Counters**

Name	Description
<i>virtual_machine</i>	<p>Amount of CPU actively being used by each virtual machine on the host. 100% represents all CPUs.</p> <p>For example, if a virtual machine has one virtual CPU that is running on a host with four CPUs and the CPU usage is 100%, the virtual machine is using one CPU resource.</p> <p>virtual CPU usage = usagemhz ÷ (number of virtual CPUs × core frequency)</p> <hr/> <p><b>Note</b> The host's view of the CPU usage, not the guest operating system view.</p> <hr/> <ul style="list-style-type: none"> <li>▪ Counter: usage</li> <li>▪ Stats Type: Rate</li> <li>▪ Unit: Percentage (%). Precision is to 1/100%. A value between 0 and 100.</li> <li>▪ Rollup Type: Average (Minimum/Maximum)</li> <li>▪ Collection Level: 1 (4)</li> </ul>

## Chart Analysis

A short spike in CPU usage or CPU ready indicates that you are making the best use of the virtual machine resources. However, if the CPU usage value for a virtual machine is above 90% and the CPU ready value is above 20%, performance is being impacted.

If performance is impacted, consider taking the following actions.

**Table 1-49. CPU Performance Enhancement Advice**

#	Resolution
1	Verify that VMware Tools is installed on every virtual machine on the host.
2	Set the CPU reservations for all high-priority virtual machines to guarantee that they receive the CPU cycles required.
3	Compare the CPU usage value of a virtual machine with the CPU usage of other virtual machines on the host or in the resource pool. The stacked line chart on the host's <b>Virtual Machine</b> view shows the CPU usage for virtual machines on the host.
4	Determine whether the high ready time for the virtual machine resulted from its CPU usage time reaching the CPU limit setting. If so, increase the CPU limit on the virtual machine.

**Table 1-49. CPU Performance Enhancement Advice (Continued)**

#	Resolution
5	Increase the CPU shares to give the virtual machine more opportunities to run. The total ready time on the host might remain at the same level if the host system is constrained by CPU. If the host ready time doesn't decrease, set the CPU reservations for high-priority virtual machines to guarantee that they receive the required CPU cycles.
6	Increase the amount of memory allocated to the virtual machine. This decreases disk and or network activity for applications that cache. This might lower disk I/O and reduce the need for the host to virtualize the hardware. Virtual machines with smaller resource allocations generally accumulate more CPU ready time.
7	Reduce the number of virtual CPUs on a virtual machine to only the number required to execute the workload. For example, a single-threaded application on a four-way virtual machine only benefits from a single vCPU. But the hypervisor's maintenance of the three idle vCPUs takes CPU cycles that could be used for other work.
8	If the host is not already in a DRS cluster, add it to one. If the host is in a DRS cluster, increase the number of hosts and migrate one or more virtual machines onto the new host.
9	Upgrade the physical CPUs or cores on the host if necessary.
1	Use the newest version of hypervisor software, and enable CPU-saving features such as TCP Segmentation Offload, large memory pages, and jumbo frames.
0	

## Disk (KBps)

The Disk (KBps) chart displays disk I/O of the host.

This chart is located in the Home view of the host **Performance** tab.

**Table 1-50. Data Counters**

Chart Label	Description
Usage	<p>Average data I/O rate across all LUNs on the host.</p> <ul style="list-style-type: none"> <li>■ Counter: usage</li> <li>■ Stats Type: Rate</li> <li>■ Unit: Kilobytes per second (KBps)</li> <li>■ Rollup Type: Average (Minimum/Maximum)</li> <li>■ Collection Level: 1 (4)</li> </ul>

## Chart Analysis

Use the disk charts to monitor average disk loads and to determine trends in disk usage. For example, you might notice a performance degradation with applications that frequently read from and write to the hard disk. If you see a spike in the number of disk read or write requests, check whether any such applications were running then.

The best ways to determine if your vSphere environment is experiencing disk problems is to monitor the disk latency data counters. You can use the advanced performance charts to view these statistics.

- The `kernellatency` data counter measures the average amount of time, in milliseconds, that the VMkernel spends processing each SCSI command. For best performance, the value must be 0 -1 milliseconds. If the value is greater than 4 ms, the virtual machines on the host are trying to send more throughput to the storage system than the configuration supports. Check the CPU usage, and increase the queue depth.

- The deviceLatency data counter measures the average amount of time, in milliseconds, to complete a SCSI command from the physical device. Depending on your hardware, a number greater than 15 ms indicates probable problems with the storage array. Move the active VMDK to a volume with more spindles or add disks to the LUN.
- The queueLatency data counter measures the average amount of time taken per SCSI command in the VMkernel queue. This value must always be zero. If not, the workload is too high and the array cannot process the data fast enough.

If the disk latency values are high, or if you notice other problems with disk I/O performance, consider taking the following actions.

**Table 1-51. Disk I/O Performance Enhancement Advice**

#	Resolution
1	<p>Increase the virtual machine memory. It allows more operating system caching, which reduces I/O activity. Note: It might require you to increase the host memory. Increasing memory might reduce the need to store data because databases can utilize the system memory to cache data and avoid disk access.</p> <p>To verify that virtual machines have adequate memory, check swap statistics in the guest operating system. Increase the guest memory, but not to an extent that leads to excessive host memory swapping. Install VMware Tools so that memory ballooning can occur.</p>
2	Defragment the file systems on all guests.
3	Disable antivirus on-demand scans on the VMDK and VMEM files.
4	Use the vendor's array tools to determine the array performance statistics. When too many servers simultaneously access common elements on an array, the disks might have trouble keeping up. To increase throughput, consider array-side improvements.
5	Use Storage vMotion to migrate I/O-intensive virtual machines across multiple hosts.
6	Balance the disk load across all physical resources available. Spread heavily used storage across LUNs that are accessed by different adapters. Use separate queues for each adapter to improve disk efficiency.
7	Configure the HBAs and RAID controllers for optimal use. Verify that the queue depths and cache settings on the RAID controllers are adequate. If not, increase the number of outstanding disk requests for the virtual machine by adjusting the <code>Disk.SchedNumReqOutstanding</code> parameter. For more information, see <i>vSphere Storage</i> .
8	For resource-intensive virtual machines, separate the virtual machine's physical disk drive from the drive with the system page file. It alleviates disk spindle contention during periods of high use.
9	On systems with sizable RAM, disable memory trimming by adding the line <code>MemTrimRate=0</code> to the virtual machine's VMX file.
10	If the combined disk I/O is higher than a single HBA capacity, use multipathing or multiple links.
11	For ESXi hosts, create virtual disks as preallocated. When you create a virtual disk for a guest operating system, select <b>Allocate all disk space now</b> . The performance degradation associated with reassigning additional disk space does not occur, and the disk is less likely to become fragmented.
12	Use the most current host hardware.

## Disk Rate (KBps)

The Disk Rate chart displays disk read and write rates for LUNs on a host, including average rates.

This chart is located in the **Home** view of the host **Performance** tab.

**Table 1-52. Data Counters**

Chart Label	Description
Read	<p>Number of disk read commands completed on each disk on the host, per second. The aggregate number of all disk read commands is also displayed in the chart.</p> <p>Read rate = blocksRead per second × blockSize</p> <ul style="list-style-type: none"> <li>■ Counter: read</li> <li>■ Stats Type: Rate</li> <li>■ Unit: Kilobytes per second (KBps)</li> <li>■ Rollup Type: Average</li> <li>■ Collection Level: 3</li> </ul>
Write	<p>Number of disk write commands completed on each disk on the host, per second. The aggregate number of all disk write commands is also displayed in the chart.</p> <p>Write rate = blocksWritten per second × blockSize</p> <ul style="list-style-type: none"> <li>■ Counter: write</li> <li>■ Stats Type: Rate</li> <li>■ Unit: Kilobytes per second (KBps)</li> <li>■ Rollup Type: Average</li> <li>■ Collection Level: 3</li> </ul>

### Chart Analysis

Use the disk charts to monitor average disk loads and to determine trends in disk usage. For example, you might notice a performance degradation with applications that frequently read from and write to the hard disk. If you see a spike in the number of disk read or write requests, check whether any such applications were running then.

The best ways to determine if your vSphere environment is experiencing disk problems is to monitor the disk latency data counters. You can use the advanced performance charts to view these statistics.

- The kernelLatency data counter measures the average amount of time, in milliseconds, that the VMkernel spends processing each SCSI command. For best performance, the value must be 0 -1 milliseconds. If the value is greater than 4 ms, the virtual machines on the host are trying to send more throughput to the storage system than the configuration supports. Check the CPU usage, and increase the queue depth.
- The deviceLatency data counter measures the average amount of time, in milliseconds, to complete a SCSI command from the physical device. Depending on your hardware, a number greater than 15 ms indicates probable problems with the storage array. Move the active VMDK to a volume with more spindles or add disks to the LUN.
- The queueLatency data counter measures the average amount of time taken per SCSI command in the VMkernel queue. This value must always be zero. If not, the workload is too high and the array cannot process the data fast enough.

If the disk latency values are high, or if you notice other problems with disk I/O performance, consider taking the following actions.

**Table 1-53. Disk I/O Performance Enhancement Advice**

#	Resolution
1	<p>Increase the virtual machine memory. It allows more operating system caching, which reduces I/O activity. Note: It might require you to increase the host memory. Increasing memory might reduce the need to store data because databases can utilize the system memory to cache data and avoid disk access.</p> <p>To verify that virtual machines have adequate memory, check swap statistics in the guest operating system. Increase the guest memory, but not to an extent that leads to excessive host memory swapping. Install VMware Tools so that memory ballooning can occur.</p>
2	Defragment the file systems on all guests.
3	Disable antivirus on-demand scans on the VMDK and VMEM files.
4	Use the vendor's array tools to determine the array performance statistics. When too many servers simultaneously access common elements on an array, the disks might have trouble keeping up. To increase throughput, consider array-side improvements.
5	Use Storage vMotion to migrate I/O-intensive virtual machines across multiple hosts.
6	Balance the disk load across all physical resources available. Spread heavily used storage across LUNs that are accessed by different adapters. Use separate queues for each adapter to improve disk efficiency.
7	Configure the HBAs and RAID controllers for optimal use. Verify that the queue depths and cache settings on the RAID controllers are adequate. If not, increase the number of outstanding disk requests for the virtual machine by adjusting the <code>Disk.SchedNumReqOutstanding</code> parameter. For more information, see <i>vSphere Storage</i> .
8	For resource-intensive virtual machines, separate the virtual machine's physical disk drive from the drive with the system page file. It alleviates disk spindle contention during periods of high use.
9	On systems with sizable RAM, disable memory trimming by adding the line <code>MemTrimRate=0</code> to the virtual machine's VMX file.
10	If the combined disk I/O is higher than a single HBA capacity, use multipathing or multiple links.
11	For ESXi hosts, create virtual disks as preallocated. When you create a virtual disk for a guest operating system, select <b>Allocate all disk space now</b> . The performance degradation associated with reassigning additional disk space does not occur, and the disk is less likely to become fragmented.
12	Use the most current host hardware.

## Disk Requests (Number)

The Disk Requests chart displays disk usage for the host.

This chart is located in the **Home** view of the host **Performance** tab.

**Table 1-54. Data Counters**

Chart Label	Description
Read Requests	<p>Number of disk read commands completed on each LUN on the host. The aggregate number of all disk read commands is also displayed in the chart.</p> <ul style="list-style-type: none"> <li>■ Counter: numberRead</li> <li>■ Stats Type: Absolute</li> <li>■ Unit: Number</li> <li>■ Rollup Type: Summation</li> <li>■ Collection Level: 3</li> </ul>
Write Requests	<p>Number of disk write commands completed on each LUN on the host. The aggregate number of all disk write commands is also displayed in the chart.</p> <ul style="list-style-type: none"> <li>■ Counter: numberWrite</li> <li>■ Stats Type: Absolute</li> <li>■ Unit: Number</li> <li>■ Rollup Type: Summation</li> <li>■ Collection Level: 3</li> </ul>

### Chart Analysis

Use the disk charts to monitor average disk loads and to determine trends in disk usage. For example, you might notice a performance degradation with applications that frequently read from and write to the hard disk. If you see a spike in the number of disk read or write requests, check whether any such applications were running then.

The best ways to determine if your vSphere environment is experiencing disk problems is to monitor the disk latency data counters. You can use the advanced performance charts to view these statistics.

- The kernelLatency data counter measures the average amount of time, in milliseconds, that the VMkernel spends processing each SCSI command. For best performance, the value must be 0 -1 milliseconds. If the value is greater than 4 ms, the virtual machines on the host are trying to send more throughput to the storage system than the configuration supports. Check the CPU usage, and increase the queue depth.
- The deviceLatency data counter measures the average amount of time, in milliseconds, to complete a SCSI command from the physical device. Depending on your hardware, a number greater than 15 ms indicates probable problems with the storage array. Move the active VMDK to a volume with more spindles or add disks to the LUN.
- The queueLatency data counter measures the average amount of time taken per SCSI command in the VMkernel queue. This value must always be zero. If not, the workload is too high and the array cannot process the data fast enough.

If the disk latency values are high, or if you notice other problems with disk I/O performance, consider taking the following actions.

**Table 1-55. Disk I/O Performance Enhancement Advice**

#	Resolution
1	<p>Increase the virtual machine memory. It allows more operating system caching, which reduces I/O activity. Note: It might require you to increase the host memory. Increasing memory might reduce the need to store data because databases can utilize the system memory to cache data and avoid disk access.</p> <p>To verify that virtual machines have adequate memory, check swap statistics in the guest operating system. Increase the guest memory, but not to an extent that leads to excessive host memory swapping. Install VMware Tools so that memory ballooning can occur.</p>
2	Defragment the file systems on all guests.
3	Disable antivirus on-demand scans on the VMDK and VMEM files.
4	Use the vendor's array tools to determine the array performance statistics. When too many servers simultaneously access common elements on an array, the disks might have trouble keeping up. To increase throughput, consider array-side improvements.
5	Use Storage vMotion to migrate I/O-intensive virtual machines across multiple hosts.
6	Balance the disk load across all physical resources available. Spread heavily used storage across LUNs that are accessed by different adapters. Use separate queues for each adapter to improve disk efficiency.
7	Configure the HBAs and RAID controllers for optimal use. Verify that the queue depths and cache settings on the RAID controllers are adequate. If not, increase the number of outstanding disk requests for the virtual machine by adjusting the <code>Disk.SchedNumReqOutstanding</code> parameter. For more information, see <i>vSphere Storage</i> .
8	For resource-intensive virtual machines, separate the virtual machine's physical disk drive from the drive with the system page file. It alleviates disk spindle contention during periods of high use.
9	On systems with sizable RAM, disable memory trimming by adding the line <code>MemTrimRate=0</code> to the virtual machine's VMX file.
10	If the combined disk I/O is higher than a single HBA capacity, use multipathing or multiple links.
11	For ESXi hosts, create virtual disks as preallocated. When you create a virtual disk for a guest operating system, select <b>Allocate all disk space now</b> . The performance degradation associated with reassigning additional disk space does not occur, and the disk is less likely to become fragmented.
12	Use the most current host hardware.

## Disk (Number)

The Disk (Number) chart displays maximum queue depth for the top ten LUNs on a host.

This chart is located in the **Home** view of the host **Performance** tab.

**Table 1-56. Data Counters**

Chart Label	Description
Maximum Queue Depth	<p>Maximum queue depth. Queue depth is the number of commands the SCSI driver queues to the HBA.</p> <ul style="list-style-type: none"> <li>■ Counter: maxQueueDepth</li> <li>■ Stats Type: Absolute</li> <li>■ Unit: Number</li> <li>■ Rollup Type: Average</li> <li>■ Collection Level: 1</li> </ul>

## Chart Analysis

Use the disk charts to monitor average disk loads and to determine trends in disk usage. For example, you might notice a performance degradation with applications that frequently read from and write to the hard disk. If you see a spike in the number of disk read or write requests, check whether any such applications were running then.

The best ways to determine if your vSphere environment is experiencing disk problems is to monitor the disk latency data counters. You can use the advanced performance charts to view these statistics.

- The `kernelLatency` data counter measures the average amount of time, in milliseconds, that the VMkernel spends processing each SCSI command. For best performance, the value must be 0 -1 milliseconds. If the value is greater than 4 ms, the virtual machines on the host are trying to send more throughput to the storage system than the configuration supports. Check the CPU usage, and increase the queue depth.
- The `deviceLatency` data counter measures the average amount of time, in milliseconds, to complete a SCSI command from the physical device. Depending on your hardware, a number greater than 15 ms indicates probable problems with the storage array. Move the active VMDK to a volume with more spindles or add disks to the LUN.
- The `queueLatency` data counter measures the average amount of time taken per SCSI command in the VMkernel queue. This value must always be zero. If not, the workload is too high and the array cannot process the data fast enough.

If the disk latency values are high, or if you notice other problems with disk I/O performance, consider taking the following actions.

**Table 1-57. Disk I/O Performance Enhancement Advice**

#	Resolution
1	<p>Increase the virtual machine memory. It allows more operating system caching, which reduces I/O activity. Note: It might require you to increase the host memory. Increasing memory might reduce the need to store data because databases can utilize the system memory to cache data and avoid disk access.</p> <p>To verify that virtual machines have adequate memory, check swap statistics in the guest operating system. Increase the guest memory, but not to an extent that leads to excessive host memory swapping. Install VMware Tools so that memory ballooning can occur.</p>
2	Defragment the file systems on all guests.
3	Disable antivirus on-demand scans on the VMDK and VMEM files.
4	Use the vendor's array tools to determine the array performance statistics. When too many servers simultaneously access common elements on an array, the disks might have trouble keeping up. To increase throughput, consider array-side improvements.
5	Use Storage vMotion to migrate I/O-intensive virtual machines across multiple hosts.
6	Balance the disk load across all physical resources available. Spread heavily used storage across LUNs that are accessed by different adapters. Use separate queues for each adapter to improve disk efficiency.
7	Configure the HBAs and RAID controllers for optimal use. Verify that the queue depths and cache settings on the RAID controllers are adequate. If not, increase the number of outstanding disk requests for the virtual machine by adjusting the <code>Disk.SchedNumReqOutstanding</code> parameter. For more information, see <i>vSphere Storage</i> .



**Table 1-57. Disk I/O Performance Enhancement Advice (Continued)**

#	Resolution
8	For resource-intensive virtual machines, separate the virtual machine's physical disk drive from the drive with the system page file. It alleviates disk spindle contention during periods of high use.
9	On systems with sizable RAM, disable memory trimming by adding the line <code>MemTrimRate=0</code> to the virtual machine's VMX file.
10	If the combined disk I/O is higher than a single HBA capacity, use multipathing or multiple links.
11	For ESXi hosts, create virtual disks as preallocated. When you create a virtual disk for a guest operating system, select <b>Allocate all disk space now</b> . The performance degradation associated with reassigning additional disk space does not occur, and the disk is less likely to become fragmented.
12	Use the most current host hardware.

## Disk (ms)

The Disk (ms) chart displays the amount of time taken to process commands on a host.

This chart is located in the **Home** view of the host **Performance** tab.

**Table 1-58. Data Counters**

Chart Label	Description
Highest Disk Latency	<p>Highest latency value of all disks used by the host.</p> <p>Latency measures the time used to process a SCSI command issued by the guest OS to the virtual machine. The kernel latency is the time VMkernel takes to process an I/O request. The device latency is the time it takes the hardware to handle the request.</p> <p>Total latency = kernelLatency + deviceLatency.</p> <ul style="list-style-type: none"> <li>■ Counter: maxTotalLatency</li> <li>■ Stats Type: Absolute</li> <li>■ Unit: Milliseconds (ms)</li> <li>■ Rollup Type: Latest (Minimum/Maximum)</li> <li>■ Collection Level: 1 (4)</li> </ul>

## Chart Analysis

Use the disk charts to monitor average disk loads and to determine trends in disk usage. For example, you might notice a performance degradation with applications that frequently read from and write to the hard disk. If you see a spike in the number of disk read or write requests, check whether any such applications were running then.

The best ways to determine if your vSphere environment is experiencing disk problems is to monitor the disk latency data counters. You can use the advanced performance charts to view these statistics.

- The kernelLatency data counter measures the average amount of time, in milliseconds, that the VMkernel spends processing each SCSI command. For best performance, the value must be 0 -1 milliseconds. If the value is greater than 4 ms, the virtual machines on the host are trying to send more throughput to the storage system than the configuration supports. Check the CPU usage, and increase the queue depth.

- The deviceLatency data counter measures the average amount of time, in milliseconds, to complete a SCSI command from the physical device. Depending on your hardware, a number greater than 15 ms indicates probable problems with the storage array. Move the active VMDK to a volume with more spindles or add disks to the LUN.
- The queueLatency data counter measures the average amount of time taken per SCSI command in the VMkernel queue. This value must always be zero. If not, the workload is too high and the array cannot process the data fast enough.

If the disk latency values are high, or if you notice other problems with disk I/O performance, consider taking the following actions.

**Table 1-59. Disk I/O Performance Enhancement Advice**

#	Resolution
1	<p>Increase the virtual machine memory. It allows more operating system caching, which reduces I/O activity. Note: It might require you to increase the host memory. Increasing memory might reduce the need to store data because databases can utilize the system memory to cache data and avoid disk access.</p> <p>To verify that virtual machines have adequate memory, check swap statistics in the guest operating system. Increase the guest memory, but not to an extent that leads to excessive host memory swapping. Install VMware Tools so that memory ballooning can occur.</p>
2	Defragment the file systems on all guests.
3	Disable antivirus on-demand scans on the VMDK and VMEM files.
4	Use the vendor's array tools to determine the array performance statistics. When too many servers simultaneously access common elements on an array, the disks might have trouble keeping up. To increase throughput, consider array-side improvements.
5	Use Storage vMotion to migrate I/O-intensive virtual machines across multiple hosts.
6	Balance the disk load across all physical resources available. Spread heavily used storage across LUNs that are accessed by different adapters. Use separate queues for each adapter to improve disk efficiency.
7	Configure the HBAs and RAID controllers for optimal use. Verify that the queue depths and cache settings on the RAID controllers are adequate. If not, increase the number of outstanding disk requests for the virtual machine by adjusting the <code>Disk.SchedNumReqOutstanding</code> parameter. For more information, see <i>vSphere Storage</i> .
8	For resource-intensive virtual machines, separate the virtual machine's physical disk drive from the drive with the system page file. It alleviates disk spindle contention during periods of high use.
9	On systems with sizable RAM, disable memory trimming by adding the line <code>MemTrimRate=0</code> to the virtual machine's VMX file.
10	If the combined disk I/O is higher than a single HBA capacity, use multipathing or multiple links.
11	For ESXi hosts, create virtual disks as preallocated. When you create a virtual disk for a guest operating system, select <b>Allocate all disk space now</b> . The performance degradation associated with reassigning additional disk space does not occur, and the disk is less likely to become fragmented.
12	Use the most current host hardware.

## Disk (KBps)

The Disk (KBps) chart displays disk usage for the 10 virtual machines on the host with the most disk usage.

This chart is located in the **Virtual Machines** view of the host **Performance** tab.

**Table 1-60. Data Counters**

Chart Label	Description
<i>virtual_machine</i>	<p>Sum of the data read from the virtual machine.</p> <ul style="list-style-type: none"> <li>■ Counter: usage</li> <li>■ Stats Type: Rate</li> <li>■ Unit: KiloBytes per second (KBps)</li> <li>■ Rollup Type: Average (Minimum/Maximum)</li> <li>■ Collection Level: 1 (4)</li> </ul>

### Chart Analysis

Use the disk charts to monitor average disk loads and to determine trends in disk usage. For example, you might notice a performance degradation with applications that frequently read from and write to the hard disk. If you see a spike in the number of disk read or write requests, check whether any such applications were running then.

The best ways to determine if your vSphere environment is experiencing disk problems is to monitor the disk latency data counters. You can use the advanced performance charts to view these statistics.

- The `kernelLatency` data counter measures the average amount of time, in milliseconds, that the VMkernel spends processing each SCSI command. For best performance, the value must be 0 -1 milliseconds. If the value is greater than 4 ms, the virtual machines on the host are trying to send more throughput to the storage system than the configuration supports. Check the CPU usage, and increase the queue depth.
- The `deviceLatency` data counter measures the average amount of time, in milliseconds, to complete a SCSI command from the physical device. Depending on your hardware, a number greater than 15 ms indicates probable problems with the storage array. Move the active VMDK to a volume with more spindles or add disks to the LUN.
- The `queueLatency` data counter measures the average amount of time taken per SCSI command in the VMkernel queue. This value must always be zero. If not, the workload is too high and the array cannot process the data fast enough.

If the disk latency values are high, or if you notice other problems with disk I/O performance, consider taking the following actions.

**Table 1-61. Disk I/O Performance Enhancement Advice**

#	Resolution
1	<p>Increase the virtual machine memory. It allows more operating system caching, which reduces I/O activity. Note: It might require you to increase the host memory. Increasing memory might reduce the need to store data because databases can utilize the system memory to cache data and avoid disk access.</p> <p>To verify that virtual machines have adequate memory, check swap statistics in the guest operating system. Increase the guest memory, but not to an extent that leads to excessive host memory swapping. Install VMware Tools so that memory ballooning can occur.</p>
2	Defragment the file systems on all guests.

**Table 1-61. Disk I/O Performance Enhancement Advice (Continued)**

#	Resolution
3	Disable antivirus on-demand scans on the VMDK and VMEM files.
4	Use the vendor's array tools to determine the array performance statistics. When too many servers simultaneously access common elements on an array, the disks might have trouble keeping up. To increase throughput, consider array-side improvements.
5	Use Storage vMotion to migrate I/O-intensive virtual machines across multiple hosts.
6	Balance the disk load across all physical resources available. Spread heavily used storage across LUNs that are accessed by different adapters. Use separate queues for each adapter to improve disk efficiency.
7	Configure the HBAs and RAID controllers for optimal use. Verify that the queue depths and cache settings on the RAID controllers are adequate. If not, increase the number of outstanding disk requests for the virtual machine by adjusting the <code>Disk.SchedNumReqOutstanding</code> parameter. For more information, see <i>vSphere Storage</i> .
8	For resource-intensive virtual machines, separate the virtual machine's physical disk drive from the drive with the system page file. It alleviates disk spindle contention during periods of high use.
9	On systems with sizable RAM, disable memory trimming by adding the line <code>MemTrimRate=0</code> to the virtual machine's VMX file.
10	If the combined disk I/O is higher than a single HBA capacity, use multipathing or multiple links.
11	For ESXi hosts, create virtual disks as preallocated. When you create a virtual disk for a guest operating system, select <b>Allocate all disk space now</b> . The performance degradation associated with reassigning additional disk space does not occur, and the disk is less likely to become fragmented.
12	Use the most current host hardware.

## Memory (%)

The Memory (%) chart displays host memory usage.

This chart is located in the **Home** view of the host **Performance** tab.

### Chart Analysis

To ensure best performance, the host memory must be large enough to accommodate the active memory of the virtual machines. The active memory can be smaller than the virtual machine memory size. It allows you to over-provision memory, but still ensures that the virtual machine active memory is smaller than the host memory.

Transient high-usage values usually do not cause performance degradation. For example, memory usage can be high when several virtual machines are started at the same time or when a spike occurs in virtual machine workload. However, a consistently high memory usage value (94% or greater) indicates that the host is probably lacking the memory required to meet the demand. If the active memory size is the same as the granted memory size, the demand for memory is greater than the memory resources available. If the active memory is consistently low, the memory size might be too large.

If the memory usage value is high, and the host has high ballooning or swapping, check the amount of free physical memory on the host. A free memory value of 6% or less indicates that the host cannot handle the demand for memory. It leads to memory reclamation, which might degrade performance.

If the host has enough free memory, check the resource shares, reservation, and limit settings of the virtual machines and resource pools on the host. Verify that the host settings are adequate and not lower than those set for the virtual machines.

If the host has little free memory available, or if you notice a degradation in performance, consider taking the following actions.

**Table 1-62. Memory Performance Enhancement Advice**

#	Resolution
1	Verify that VMware Tools is installed on each virtual machine. The balloon driver is installed with VMware Tools and is critical to performance.
2	Verify that the balloon driver is enabled. The VMkernel regularly reclaims unused virtual machine memory by ballooning and swapping. Generally, it does not impact virtual machine performance.
3	Reduce the memory space on the virtual machine, and correct the cache size if it is too large. This frees up memory for other virtual machines.
4	If the memory reservation of the virtual machine is set to a value much higher than its active memory, decrease the reservation setting so that the VMkernel can reclaim the idle memory for other virtual machines on the host.
5	Migrate one or more virtual machines to a host in a DRS cluster.
6	Add physical memory to the host.

## Memory (Balloon)

The Memory (Balloon) chart displays balloon memory on a host.

This chart is located in the **Home** view of the host **Performance** tab.

**Table 1-63. Data Counters**

Chart Label	Description
Balloon	<p>Sum of the guest physical memory reclaimed by the balloon driver for all powered on virtual machines on the host.</p> <ul style="list-style-type: none"> <li>■ Counter: vmmemctl</li> <li>■ Stats Type: Absolute</li> <li>■ Unit: Megabytes (MB)</li> <li>■ Rollup Type: Average (Minimum/Maximum)</li> <li>■ Collection Level: 1 (4)</li> </ul>

## Chart Analysis

To ensure best performance, the host memory must be large enough to accommodate the active memory of the virtual machines. The active memory can be smaller than the virtual machine memory size. It allows you to over-provision memory, but still ensures that the virtual machine active memory is smaller than the host memory.

Transient high-usage values usually do not cause performance degradation. For example, memory usage can be high when several virtual machines are started at the same time or when a spike occurs in virtual machine workload. However, a consistently high memory usage value (94% or greater) indicates that the host is probably lacking the memory required to meet the demand. If the active memory size is the same as the granted memory size, the demand for memory is greater than the memory resources available. If the active memory is consistently low, the memory size might be too large.

If the memory usage value is high, and the host has high ballooning or swapping, check the amount of free physical memory on the host. A free memory value of 6% or less indicates that the host cannot handle the demand for memory. It leads to memory reclamation, which might degrade performance.

If the host has enough free memory, check the resource shares, reservation, and limit settings of the virtual machines and resource pools on the host. Verify that the host settings are adequate and not lower than those set for the virtual machines.

If the host has little free memory available, or if you notice a degradation in performance, consider taking the following actions.

**Table 1-64. Memory Performance Enhancement Advice**

#	Resolution
1	Verify that VMware Tools is installed on each virtual machine. The balloon driver is installed with VMware Tools and is critical to performance.
2	Verify that the balloon driver is enabled. The VMkernel regularly reclaims unused virtual machine memory by ballooning and swapping. Generally, it does not impact virtual machine performance.
3	Reduce the memory space on the virtual machine, and correct the cache size if it is too large. This frees up memory for other virtual machines.
4	If the memory reservation of the virtual machine is set to a value much higher than its active memory, decrease the reservation setting so that the VMkernel can reclaim the idle memory for other virtual machines on the host.
5	Migrate one or more virtual machines to a host in a DRS cluster.
6	Add physical memory to the host.

## Memory (MBps)

The Memory (MBps) chart displays the swap in and swap out rates for a host.

This chart is located on the **Home** view of the Host **Performance** tab.

**Table 1-65. Data Counters**

Chart Label	Description
swapinRate	<p>Average rate at which memory is swapped in from the host swap file.</p> <ul style="list-style-type: none"> <li>■ Counter: swapinRate</li> <li>■ Stats Type: Rate</li> <li>■ Unit: MegaBytes per second (MBps)</li> <li>■ Rollup Type: Average (Minimum/Maximum)</li> <li>■ Collection Level: 1 (4)</li> </ul>
swapoutRate	<p>Average rate at which memory is swapped out to the host swap file.</p> <ul style="list-style-type: none"> <li>■ Counter: swapoutRate</li> <li>■ Stats Type: Rate</li> <li>■ Unit: MegaBytes per second (MBps)</li> <li>■ Rollup Type: Average (Minimum/Maximum)</li> <li>■ Collection Level: 1 (4)</li> </ul>

### Chart Analysis

Host memory must be large enough to accommodate virtual machine workload. Transient high-usage values usually do not cause performance degradation. For example, memory usage can be high when several virtual machines are started at the same time or when there is a spike in virtual machine workload.

However, a consistently high memory usage value (94% or greater) indicates the host does not have the memory resources required to meet the demand. If the memory balloon and swap values are not high, performance is probably not affected. If the memory usage value is high, and the host has high ballooning or swapping, check the amount of free physical memory on the host. A free memory value of 6% or less indicates that the host requires more memory resources.

If the host is not lacking memory resources, check the resource shares, reservation, and limit settings of the virtual machines and resource pools on the host. Verify that the host settings are adequate and not lower than those set for the virtual machines.

If the host is lacking memory resources or you notice a degradation in performance, consider taking the following actions.

**Table 1-66. Memory Performance Enhancement Advice**

#	Resolution
1	Verify that VMware Tools is installed on each virtual machine. The balloon driver is installed with VMware Tools and is critical to performance.
2	Verify that the balloon driver is enabled. The VMkernel regularly reclaims unused virtual machine memory by ballooning and swapping. Generally, this does not impact virtual machine performance.
3	Reduce the memory space on the virtual machine, and correct the cache size if it is too large. This frees up memory for other virtual machines.
4	If the memory reservation of a virtual machine is set to a value much higher than its active memory, decrease the reservation setting so that the VMkernel can reclaim the idle memory for other virtual machines on the host.

**Table 1-66. Memory Performance Enhancement Advice (Continued)**

#	Resolution
5	Migrate one or more virtual machines to a host in a DRS cluster.
6	Add physical memory to the host.

## Memory (MB)

The Memory (MB) chart displays memory data counters for hosts.

This chart is located in the **Home** view of the host **Performance** tab.

**Note** Guest physical memory refers to the virtual hardware memory presented to a virtual machine for its guest operating system. Machine memory is the actual physical RAM in the host.

Not all counters are collected at collection level 1.

**Table 1-67. Data Counters**

Chart Label	Description
Active	<p>Sum of the active guest physical memory of all powered on virtual machines on the host, plus memory used by basic VMkernel applications. Active memory is estimated by the VMkernel and is based on the current workload of the host.</p> <ul style="list-style-type: none"> <li>■ Counter: active</li> <li>■ Stats Type: Absolute</li> <li>■ Unit: Megabytes (MB)</li> <li>■ Rollup Type: Average (Minimum/Maximum)</li> <li>■ Collection Level: 2 (4)</li> </ul>
Balloon	<p>Sum of the guest physical memory reclaimed by the balloon driver for all powered on virtual machines on the host.</p> <ul style="list-style-type: none"> <li>■ Counter: vmmemctl</li> <li>■ Stats Type: Absolute</li> <li>■ Unit: Megabytes (MB)</li> <li>■ Rollup Type: Average (Minimum/Maximum)</li> <li>■ Collection Level: 1 (4)</li> </ul>
Balloon Target	<p>Sum of the balloon target memory of all powered on virtual machines on the host.</p> <p>If the balloon target value is greater than the balloon value, the VMkernel inflates the balloon, causing more virtual machine memory to be reclaimed. If the balloon target value is less than the balloon value, the VMkernel deflates the balloon, which allows the virtual machine to consume additional memory if needed.</p> <p>Virtual machines initiate memory reallocation. Therefore, it is possible to have a balloon target value of 0 and a balloon value greater than 0.</p> <ul style="list-style-type: none"> <li>■ Counter: vmmemctltarget</li> <li>■ Stats Type: Absolute</li> <li>■ Unit: Megabytes (MB)</li> <li>■ Rollup Type: Average (Minimum/Maximum)</li> <li>■ Collection Level: 2 (4)</li> </ul>



**Table 1-67. Data Counters (Continued)**

Chart Label	Description
Consumed	<p>Amount of machine memory used on the host.</p> <p>Consumed memory includes virtual machine memory, service console memory, and VMkernel memory.</p> <p>consumed memory = total host memory - free host memory</p> <ul style="list-style-type: none"> <li>■ Counter: consumed</li> <li>■ Stats Type: Absolute</li> <li>■ Unit: Megabytes (MB)</li> <li>■ Rollup Type: Average (Minimum/Maximum)</li> <li>■ Collection Level: 1 (4)</li> </ul>
Granted	<p>Sum of the guest physical memory granted for all powered on virtual machines. Granted memory is mapped to the host's machine memory.</p> <p>Granted memory for a host includes the shared memory of each virtual machine on the host.</p> <ul style="list-style-type: none"> <li>■ Counter: granted</li> <li>■ Stats Type: Absolute</li> <li>■ Unit: Megabytes (MB)</li> <li>■ Rollup Type: Average (Minimum/Maximum)</li> <li>■ Collection Level: 2 (4)</li> </ul>
Shared Common	<p>Amount of machine memory shared by all powered on virtual machines.</p> <p>Shared common memory consists of the entire pool of memory from which sharing is possible, including the amount of physical RAM required by the guest memory.</p> <p>memory shared - memory shared common = amount of memory saved on the host from sharing</p> <ul style="list-style-type: none"> <li>■ Counter: sharedcommon</li> <li>■ Stats Type: Absolute</li> <li>■ Unit: Megabytes (MB)</li> <li>■ Rollup Type: Average (Minimum/Maximum)</li> <li>■ Collection Level: 2 (4)</li> </ul>
Swap Used	<p>Sum of the memory swapped by all powered on virtual machines on the host.</p> <ul style="list-style-type: none"> <li>■ Counter: swapped</li> <li>■ Stats Type: Absolute</li> <li>■ Unit: Megabytes (MB)</li> <li>■ Rollup Type: Average (Minimum/Maximum)</li> <li>■ Collection Level: 2 (4)</li> </ul>

### Chart Analysis

To ensure best performance, the host memory must be large enough to accommodate the active memory of the virtual machines. The active memory can be smaller than the virtual machine memory size. It allows you to over-provision memory, but still ensures that the virtual machine active memory is smaller than the host memory.

Transient high-usage values usually do not cause performance degradation. For example, memory usage can be high when several virtual machines are started at the same time or when a spike occurs in virtual machine workload. However, a consistently high memory usage value (94% or greater) indicates that the host is probably lacking the memory required to meet the demand. If the active memory size is the same as the granted memory size, the demand for memory is greater than the memory resources available. If the active memory is consistently low, the memory size might be too large.

If the memory usage value is high, and the host has high ballooning or swapping, check the amount of free physical memory on the host. A free memory value of 6% or less indicates that the host cannot handle the demand for memory. It leads to memory reclamation, which might degrade performance.

If the host has enough free memory, check the resource shares, reservation, and limit settings of the virtual machines and resource pools on the host. Verify that the host settings are adequate and not lower than those set for the virtual machines.

If the host has little free memory available, or if you notice a degradation in performance, consider taking the following actions.

**Table 1-68. Memory Performance Enhancement Advice**

#	Resolution
1	Verify that VMware Tools is installed on each virtual machine. The balloon driver is installed with VMware Tools and is critical to performance.
2	Verify that the balloon driver is enabled. The VMkernel regularly reclaims unused virtual machine memory by ballooning and swapping. Generally, it does not impact virtual machine performance.
3	Reduce the memory space on the virtual machine, and correct the cache size if it is too large. This frees up memory for other virtual machines.
4	If the memory reservation of the virtual machine is set to a value much higher than its active memory, decrease the reservation setting so that the VMkernel can reclaim the idle memory for other virtual machines on the host.
5	Migrate one or more virtual machines to a host in a DRS cluster.
6	Add physical memory to the host.

## Memory Usage

The Memory Usage chart displays memory usage for the 10 virtual machines on the host with the most memory usage.

This chart is located in the **Virtual Machines** view of the host **Performance** tab.

## Virtual Machine Counters

**Note** Guest physical memory refers to the virtual hardware memory presented to a virtual machine for its guest operating system.

**Table 1-69. Data Counters**

Chart Label	Description
Usage	<p>Amount of guest physical memory currently in use on the virtual machine.</p> <ul style="list-style-type: none"> <li>■ Counter: usage</li> <li>■ Stats Type: Absolute</li> <li>■ Unit: Percentage (%)</li> <li>■ Rollup Type: Average (Minimum/Maximum)</li> <li>■ Collection Level: 1 (4)</li> </ul>

### Chart Analysis

A virtual machine's memory size must be slightly larger than the average guest memory usage. This enables the host to accommodate workload spikes without swapping memory among guests. Increasing the virtual machine memory size results in more overhead memory usage.

If sufficient swap space is available, a high balloon value does not cause performance problems. However, if the `swpin` and `swpout` values for the host are large, the host is probably lacking the amount of memory required to meet the demand.

If a virtual machine has high ballooning or swapping, check the amount of free physical memory on the host. A free memory value of 6% or less indicates that the host cannot meet the memory requirements. This leads to memory reclamation, which might degrade performance. If the active memory size is the same as the granted memory size, demand for memory is greater than the memory resources available. If the active memory is consistently low, the memory size might be too large.

If the host has enough free memory, check the resource shares, reservation, and limit of the virtual machines and resource pools on the host. Verify that the host settings are adequate and not lower than those set for the virtual machine.

If little free memory is available, or if you notice degradation in performance, consider taking the following actions.

**Table 1-70. Memory Performance Enhancement Advice**

#	Resolution
1	Verify that VMware Tools is installed on each virtual machine. The balloon driver is installed with VMware Tools and is critical to performance.
2	Verify that the balloon driver is enabled. The VMkernel regularly reclaims unused virtual machine memory by ballooning and swapping. Generally, this does not impact virtual machine performance.
3	Reduce the memory space on the virtual machine, and correct the cache size if it is too large. This frees up memory for other virtual machines.
4	If the memory reservation of the virtual machine is set to a value much higher than its active memory, decrease the reservation setting so that the VMkernel can reclaim the idle memory for other virtual machines on the host.
5	Migrate one or more virtual machines to a host in a DRS cluster.
6	Add physical memory to the host.

## Network (Mbps)

The Network (Mbps) chart displays network usage for the host.

This chart is located in the **Home** view of the Host **Performance** tab.

**Table 1-71. Host Counters**

Chart Label	Description
Usage	<p>Average rate at which data is transmitted and received across all NIC instances connected to the host.</p> <ul style="list-style-type: none"> <li>■ Counter: usage</li> <li>■ Stats Type: Rate</li> <li>■ Unit: Megabits per second (Mbps)</li> <li>■ Rollup Type: Average (Minimum/Maximum)</li> <li>■ Collection Level: 1 (4)</li> </ul>

### Chart Analysis

Network performance depends on the application workload and network configuration. Dropped network packets indicate a bottleneck in the network. To determine whether packets are being dropped, use `esx top` or the advanced performance charts to examine the `droppedTx` and `droppedRx` network counter values.

If packets are being dropped, adjust the virtual machine shares. If packets are not being dropped, check the size of the network packets and the data receive and transfer rates. In general, the larger the network packets, the faster the network speed. When the packet size is large, fewer packets are transferred, which reduces the amount of CPU required to process the data. When network packets are small, more packets are transferred but the network speed is slower because more CPU is required to process the data.

---

**Note** In some instances, large packets might result in a high network latency. To check the network latency, use the VMware AppSpeed performance monitoring application or a third-party application.

---

If packets are not being dropped and the data receive rate is slow, the host is probably lacking the CPU resources required to handle the load. Check the number of virtual machines assigned to each physical NIC. If necessary, perform load balancing by moving virtual machines to different vSwitches or by adding more NICs to the host. You can also move virtual machines to another host or increase the host CPU or virtual machine CPU.

If you experience network-related performance problems, also consider taking the following actions.

**Table 1-72. Networking Performance Enhancement Advice**

#	Resolution
1	Verify that VMware Tools is installed on each virtual machine.
2	If possible, use vmxnet3 NIC drivers, which are available with VMware Tools. They are optimized for high performance.
3	If virtual machines running on the same host communicate with each other, connect them to the same vSwitch to avoid transferring packets over the physical network.

**Table 1-72. Networking Performance Enhancement Advice (Continued)**

#	Resolution
4	Assign each physical NIC to a port group and a vSwitch.
5	Use separate physical NICs to handle the different traffic streams, such as network packets generated by virtual machines, iSCSI protocols, vMotion tasks.
6	Ensure that the physical NIC capacity is large enough to handle the network traffic on that vSwitch. If the capacity is not enough, consider using a high-bandwidth physical NIC (10 Gbps). Alternatively, consider moving some virtual machines to a vSwitch with a lighter load or to a new vSwitch.
7	If packets are being dropped at the vSwitch port, increase the virtual network driver ring buffers where applicable.
8	Verify that the reported speed and duplex settings for the physical NIC match the hardware expectations and that the hardware is configured to run at its maximum capability. For example, verify that NICs with 1 Gbps are not reset to 100 Mbps because they are connected to an older switch.
9	Verify that all NICs are running in full duplex mode. Hardware connectivity problems might result in a NIC resetting itself to a lower speed or half duplex mode.
10	Use vNICs that are TCP Segmentation Offload (TSO)-capable, and verify that TSO-Jumbo Frames are enabled where possible.

## Network Rate (Mbps)

The Network Rate chart displays network bandwidth on a host.

The Network Data Transmitted/Received chart for hosts is located in the **Home** view of the Host **Performance** tab.

**Table 1-73. Data Counters**

Chart Label	Description
Data Receive Rate	<p>Rate at which data is received across the top ten physical NIC instances on the host. This represents the bandwidth of the network. The chart also displays the aggregated data receive rate of all physical NICs.</p> <ul style="list-style-type: none"> <li>■ Counter: received</li> <li>■ Stats Type: Rate</li> <li>■ Unit: Megabits per second (Mbps)</li> <li>■ Rollup Type: Average</li> <li>■ Collection Level: 3 (4)</li> </ul>
Data Transmit Rate	<p>Rate at which data is transmitted across the top ten physical NIC instances on the host. This represents the bandwidth of the network. The chart also displays the aggregated data transmit rate of all physical NICs.</p> <ul style="list-style-type: none"> <li>■ Counter: transmitted</li> <li>■ Stats Type: Rate</li> <li>■ Unit: Megabits per second (Mbps)</li> <li>■ Rollup Type: Average</li> <li>■ Collection Level: 3 (4)</li> </ul>

## Chart Analysis

Network performance depends on the application workload and network configuration. Dropped network packets indicate a bottleneck in the network. To determine whether packets are being dropped, use `esxtop` or the advanced performance charts to examine the `droppedTx` and `droppedRx` network counter values.

If packets are being dropped, adjust the virtual machine shares. If packets are not being dropped, check the size of the network packets and the data receive and transfer rates. In general, the larger the network packets, the faster the network speed. When the packet size is large, fewer packets are transferred, which reduces the amount of CPU required to process the data. When network packets are small, more packets are transferred but the network speed is slower because more CPU is required to process the data.

---

**Note** In some instances, large packets might result in a high network latency. To check the network latency, use the VMware AppSpeed performance monitoring application or a third-party application.

---

If packets are not being dropped and the data receive rate is slow, the host is probably lacking the CPU resources required to handle the load. Check the number of virtual machines assigned to each physical NIC. If necessary, perform load balancing by moving virtual machines to different vSwitches or by adding more NICs to the host. You can also move virtual machines to another host or increase the host CPU or virtual machine CPU.

If you experience network-related performance problems, also consider taking the following actions.

**Table 1-74. Networking Performance Enhancement Advice**

#	Resolution
1	Verify that VMware Tools is installed on each virtual machine.
2	If possible, use <code>vmxnet3</code> NIC drivers, which are available with VMware Tools. They are optimized for high performance.
3	If virtual machines running on the same host communicate with each other, connect them to the same vSwitch to avoid transferring packets over the physical network.
4	Assign each physical NIC to a port group and a vSwitch.
5	Use separate physical NICs to handle the different traffic streams, such as network packets generated by virtual machines, iSCSI protocols, vMotion tasks.
6	Ensure that the physical NIC capacity is large enough to handle the network traffic on that vSwitch. If the capacity is not enough, consider using a high-bandwidth physical NIC (10 Gbps). Alternatively, consider moving some virtual machines to a vSwitch with a lighter load or to a new vSwitch.
7	If packets are being dropped at the vSwitch port, increase the virtual network driver ring buffers where applicable.
8	Verify that the reported speed and duplex settings for the physical NIC match the hardware expectations and that the hardware is configured to run at its maximum capability. For example, verify that NICs with 1 Gbps are not reset to 100 Mbps because they are connected to an older switch.
9	Verify that all NICs are running in full duplex mode. Hardware connectivity problems might result in a NIC resetting itself to a lower speed or half duplex mode.
10	Use vNICs that are TCP Segmentation Offload (TSO)-capable, and verify that TSO-Jumbo Frames are enabled where possible.

## Network Packets (Number)

The Network Packets chart displays the network bandwidth on a host.

This chart is located in the **Home** view of the Host **Performance** tab.

**Table 1-75. Data Counters**

Chart Label	Description
Packets Received	<p>Number of network packets received across the top ten physical NIC instances on the host. The chart also displays the aggregated value for all NICs.</p> <ul style="list-style-type: none"> <li>■ Counter: packetRx</li> <li>■ Stats Type: Absolute</li> <li>■ Unit: Number</li> <li>■ Rollup Type: Summation</li> <li>■ Collection Level: 3</li> </ul>
Packets Transmitted	<p>Number of network packets transmitted across the top ten physical NIC instances on the host. The chart also displays the aggregated value for all NICs.</p> <ul style="list-style-type: none"> <li>■ Counter: packetTx</li> <li>■ Stats Type: Absolute</li> <li>■ Unit: Number</li> <li>■ Rollup Type: Summation</li> <li>■ Collection Level: 3</li> </ul>

### Chart Analysis

Network performance depends on the application workload and network configuration. Dropped network packets indicate a bottleneck in the network. To determine whether packets are being dropped, use `esx top` or the advanced performance charts to examine the `droppedTx` and `droppedRx` network counter values.

If packets are being dropped, adjust the virtual machine shares. If packets are not being dropped, check the size of the network packets and the data receive and transfer rates. In general, the larger the network packets, the faster the network speed. When the packet size is large, fewer packets are transferred, which reduces the amount of CPU required to process the data. When network packets are small, more packets are transferred but the network speed is slower because more CPU is required to process the data.

---

**Note** In some instances, large packets might result in a high network latency. To check the network latency, use the VMware AppSpeed performance monitoring application or a third-party application.

---

If packets are not being dropped and the data receive rate is slow, the host is probably lacking the CPU resources required to handle the load. Check the number of virtual machines assigned to each physical NIC. If necessary, perform load balancing by moving virtual machines to different vSwitches or by adding more NICs to the host. You can also move virtual machines to another host or increase the host CPU or virtual machine CPU.

If you experience network-related performance problems, also consider taking the following actions.

**Table 1-76. Networking Performance Enhancement Advice**

#	Resolution
1	Verify that VMware Tools is installed on each virtual machine.
2	If possible, use vmxnet3 NIC drivers, which are available with VMware Tools. They are optimized for high performance.
3	If virtual machines running on the same host communicate with each other, connect them to the same vSwitch to avoid transferring packets over the physical network.
4	Assign each physical NIC to a port group and a vSwitch.
5	Use separate physical NICs to handle the different traffic streams, such as network packets generated by virtual machines, iSCSI protocols, vMotion tasks.
6	Ensure that the physical NIC capacity is large enough to handle the network traffic on that vSwitch. If the capacity is not enough, consider using a high-bandwidth physical NIC (10 Gbps). Alternatively, consider moving some virtual machines to a vSwitch with a lighter load or to a new vSwitch.
7	If packets are being dropped at the vSwitch port, increase the virtual network driver ring buffers where applicable.
8	Verify that the reported speed and duplex settings for the physical NIC match the hardware expectations and that the hardware is configured to run at its maximum capability. For example, verify that NICs with 1 Gbps are not reset to 100 Mbps because they are connected to an older switch.
9	Verify that all NICs are running in full duplex mode. Hardware connectivity problems might result in a NIC resetting itself to a lower speed or half duplex mode.
10	Use vNICs that are TCP Segmentation Offload (TSO)-capable, and verify that TSO-Jumbo Frames are enabled where possible.

## Network (Mbps)

The Network (Mbps) chart displays network usage for the 10 virtual machines on the host with the most network usage.

This chart is located in the **Virtual Machines** view of the host **Performance** tab.

**Table 1-77. Data Counters**

Chart Label	Description
<virtual machine>	<p>Sum of the data transmitted and received across all virtual NIC instances connected to the virtual machine.</p> <ul style="list-style-type: none"> <li>■ Counter: usage</li> <li>■ Stats Type: Rate</li> <li>■ Unit: Megabits per second (Mbps)</li> <li>■ Rollup Type: Average (Minimum/Maximum)</li> <li>■ Collection Level: 1 (4)</li> </ul>

## Chart Analysis

Network performance depends on the application workload and network configuration. Dropped network packets indicate a bottleneck in the network. To determine whether packets are being dropped, use `esxtop` or the advanced performance charts to examine the `droppedTx` and `droppedRx` network counter values.



If packets are being dropped, adjust the virtual machine shares. If packets are not being dropped, check the size of the network packets and the data receive and transfer rates. In general, the larger the network packets, the faster the network speed. When the packet size is large, fewer packets are transferred, which reduces the amount of CPU required to process the data. When network packets are small, more packets are transferred but the network speed is slower because more CPU is required to process the data.

---

**Note** In some instances, large packets might result in a high network latency. To check the network latency, use the VMware AppSpeed performance monitoring application or a third-party application.

---

If packets are not being dropped and the data receive rate is slow, the host is probably lacking the CPU resources required to handle the load. Check the number of virtual machines assigned to each physical NIC. If necessary, perform load balancing by moving virtual machines to different vSwitches or by adding more NICs to the host. You can also move virtual machines to another host or increase the host CPU or virtual machine CPU.

If you experience network-related performance problems, also consider taking the following actions.

**Table 1-78. Networking Performance Enhancement Advice**

#	Resolution
1	Verify that VMware Tools is installed on each virtual machine.
2	If possible, use vmxnet3 NIC drivers, which are available with VMware Tools. They are optimized for high performance.
3	If virtual machines running on the same host communicate with each other, connect them to the same vSwitch to avoid transferring packets over the physical network.
4	Assign each physical NIC to a port group and a vSwitch.
5	Use separate physical NICs to handle the different traffic streams, such as network packets generated by virtual machines, iSCSI protocols, vMotion tasks.
6	Ensure that the physical NIC capacity is large enough to handle the network traffic on that vSwitch. If the capacity is not enough, consider using a high-bandwidth physical NIC (10 Gbps). Alternatively, consider moving some virtual machines to a vSwitch with a lighter load or to a new vSwitch.
7	If packets are being dropped at the vSwitch port, increase the virtual network driver ring buffers where applicable.
8	Verify that the reported speed and duplex settings for the physical NIC match the hardware expectations and that the hardware is configured to run at its maximum capability. For example, verify that NICs with 1 Gbps are not reset to 100 Mbps because they are connected to an older switch.
9	Verify that all NICs are running in full duplex mode. Hardware connectivity problems might result in a NIC resetting itself to a lower speed or half duplex mode.
10	Use vNICs that are TCP Segmentation Offload (TSO)-capable, and verify that TSO-Jumbo Frames are enabled where possible.

## Resource Pools

The resource pool charts contain information about CPU and memory usage for resource pools. The help topic for each chart contains information about the data counters displayed in that chart. The counters available are determined by the collection level set for vCenter Server.

## CPU (MHz)

The CPU (MHz) chart displays CPU usage in the resource pool or vApp.

This chart is located in the Home view of the Resource Pool or vApp **Performance** tab.

### Counters

**Table 1-79. Data Counters**

Chart Label	Description
Usage	<p>CPU usage is the sum of the average CPU usage values of the virtual machines in the resource pool or vApp.</p> <p>CPU usage = number of cores * CPU frequency</p> <ul style="list-style-type: none"> <li>■ Counter: usagemhz</li> <li>■ Stats Type: Rate</li> <li>■ Unit: Megahertz (MHz)</li> <li>■ Rollup Type: Average (Minimum/Maximum)</li> <li>■ Collection Level: 1 (4)</li> </ul>

### Chart Analysis

A short spike in CPU usage indicates that you are making the best use of the resources available. However, if the value is constantly high, the CPU demanded is likely greater than the CPU capacity available. A high CPU usage value can lead to increased ready time and processor queuing of the virtual machines in the resource pool. Generally, if the CPU usage value for a virtual machine is above 90% and the CPU ready value for a virtual machine is above 20%, performance is impacted.

If performance is impacted, consider taking the following actions.

**Table 1-80. CPU Performance Enhancement Advice**

#	Resolution
1	Verify that VMware Tools is installed on each virtual machine.
2	Deploy single-threaded applications on uniprocessor virtual machines instead of SMP virtual machines.
3	Migrate one or more virtual machines to a new host.
4	Upgrade the physical CPUs or cores on each host if necessary.
5	Enable CPU-saving features, such as TCP Segmentation Offload.
6	Replace software I/O with dedicated hardware, such as iSCSI HBAs or TCP Segmentation Offload NICs.

## CPU Usage

The CPU Usage chart displays CPU usage of virtual machines in the resource pool or vApp. The chart displays the top 10 virtual machines with the highest CPU usage.

This chart is located in the Resource Pools & Virtual Machines view of the Resource Pool or vApp **Performance** tab.

**Table 1-81. Data Counters**

Chart Label	Description
<i>virtual_machine</i>	<p>Amount of CPU actively used by virtual machines.</p> <ul style="list-style-type: none"> <li>■ Counter: usagemhz</li> <li>■ Stats Type: Rate</li> <li>■ Unit: Megahertz (MHz)</li> <li>■ Rollup Type: Average (Minimum/Maximum)</li> <li>■ Collection Level: 1 (4)</li> </ul>

### Chart Analysis

A short spike in CPU usage or CPU ready indicates that you are making the best use of the virtual machine resources. However, if the CPU usage value for a virtual machine is above 90% and the CPU ready value is above 20%, performance is being impacted.

If performance is impacted, consider taking the following actions.

**Table 1-82. CPU Performance Enhancement Advice**

#	Resolution
1	Verify that VMware Tools is installed on every virtual machine on the host.
2	Set the CPU reservations for all high-priority virtual machines to guarantee that they receive the CPU cycles required.
3	Compare the CPU usage value of a virtual machine with the CPU usage of other virtual machines on the host or in the resource pool. The stacked line chart on the host's <b>Virtual Machine</b> view shows the CPU usage for virtual machines on the host.
4	Determine whether the high ready time for the virtual machine resulted from its CPU usage time reaching the CPU limit setting. If so, increase the CPU limit on the virtual machine.
5	Increase the CPU shares to give the virtual machine more opportunities to run. The total ready time on the host might remain at the same level if the host system is constrained by CPU. If the host ready time doesn't decrease, set the CPU reservations for high-priority virtual machines to guarantee that they receive the required CPU cycles.
6	Increase the amount of memory allocated to the virtual machine. This decreases disk and or network activity for applications that cache. This might lower disk I/O and reduce the need for the host to virtualize the hardware. Virtual machines with smaller resource allocations generally accumulate more CPU ready time.
7	Reduce the number of virtual CPUs on a virtual machine to only the number required to execute the workload. For example, a single-threaded application on a four-way virtual machine only benefits from a single vCPU. But the hypervisor's maintenance of the three idle vCPUs takes CPU cycles that could be used for other work.
8	If the host is not already in a DRS cluster, add it to one. If the host is in a DRS cluster, increase the number of hosts and migrate one or more virtual machines onto the new host.
9	Upgrade the physical CPUs or cores on the host if necessary.
10	Use the newest version of hypervisor software, and enable CPU-saving features such as TCP Segmentation Offload, large memory pages, and jumbo frames.

### Memory (MB)

The Memory (MB) chart displays memory usage in the resource pool or vApp.

This chart is located in the **Home** view of the resource pool or vApp **Performance** tab.

**Table 1-83. Data Counters**

Chart Label	Description
<i>resource_pool</i> or <i>vApp</i>	<p>Sum of the active memory used by all virtual machines in the resource pool or vApp. Active memory is determined by the VMkernel and includes overhead memory.</p> <p>memory usage = active memory / configured virtual machine memory size</p> <ul style="list-style-type: none"> <li>■ Counter: used</li> <li>■ Stats Type: Absolute</li> <li>■ Unit: Megabytes (MB)</li> <li>■ Rollup Type: Average</li> <li>■ Collection Level: 1</li> </ul>

### Chart Analysis

Memory usage is not an indicator of performance problems. Memory can be high if a host is swapping or ballooning, which can result in virtual machine guest swapping. In such cases, check for other problems, such as CPU over-commitment or storage latencies.

If you have constantly high memory usage in a cluster, resource pool, or vApp, consider taking the following actions.

**Table 1-84. Memory Performance Enhancement Advice**

#	Resolution
1	Verify that VMware Tools is installed on each virtual machine.
2	Verify that the balloon driver is enabled. The balloon driver is installed with VMware Tools and is critical to performance. The VMkernel regularly reclaims unused virtual machine memory by ballooning and swapping. Generally, it does not impact virtual machine performance.
3	If the balloon value is high, check the resource shares, reservations, and limits for the virtual machines and resource pools on the hosts. Verify that the host's settings are adequate and not lower than those set for the virtual machine. If free memory is available on the hosts and the virtual machines are experiencing high swap or balloon memory, the virtual machine (or resource pool, if it belongs to one) has reached its resource limit. Check the maximum resource limit set on that host.
4	<p>If the cluster is not a DRS cluster, enable DRS. To enable DRS, perform the following tasks:</p> <ol style="list-style-type: none"> <li>1 Select the cluster, and click the <b>Configure</b> tab.</li> <li>2 Under <b>Services</b>, click <b>vSphere DRS</b>.</li> <li>3 click <b>Edit</b>.</li> </ol> <p>An Edit Cluster Settings dialog box opens.</p> <ol style="list-style-type: none"> <li>4 Click <b>Turn ON vSphere DRS</b>, and click <b>OK</b>.</li> </ol>
5	<p>If the cluster is a DRS cluster:</p> <ul style="list-style-type: none"> <li>■ Increase the number of hosts, and migrate one or more virtual machines to the new host.</li> <li>■ Check the aggressiveness threshold. If the value is low, increase the threshold. It might help avoid hot spots in the cluster.</li> </ul>
6	Add more physical memory to one or more hosts.

### Memory Consumed

The Memory Consumed chart displays the memory performance of all virtual machines in the resource pool or vApp.

This chart is located in the **Resource Pools & Virtual Machines** view of the resource pool or vApp **Performance** tab.

For resource pools and virtual machines in a resource pool or vApp, this chart is located in the **Resource Pools & Virtual Machines** view of the resource pool or vApp **Performance** tab.

**Table 1-85. Data Counters**

Chart Label	Description
<i>virtual_machine</i>	<p>Amount of host memory used by the virtual machine for its guest operating system's physical memory. Memory overhead is not included in consumed memory.</p> <p>consumed memory = memory granted - memory saved from page sharing</p> <p>For example, if a virtual machine has 100 MB of memory that is shared equally with three other virtual machines, its portion of the shared memory is 25 MB (100 MB ÷ 4 VMs). This amount is counted in the memory consumed data counter.</p> <ul style="list-style-type: none"> <li>■ Counter: consumed</li> <li>■ Stats Type: Absolute</li> <li>■ Unit: Megabytes (MB)</li> <li>■ Rollup Type: Average (Minimum/Maximum)</li> <li>■ Collection Level: 1 (4)</li> </ul>

### Chart Analysis

A virtual machine's memory size must be slightly larger than the average guest memory usage. This enables the host to accommodate workload spikes without swapping memory among guests. Increasing the virtual machine memory size results in more overhead memory usage.

If sufficient swap space is available, a high balloon value does not cause performance problems. However, if the `swpin` and `swpout` values for the host are large, the host is probably lacking the amount of memory required to meet the demand.

If a virtual machine has high ballooning or swapping, check the amount of free physical memory on the host. A free memory value of 6% or less indicates that the host cannot meet the memory requirements. This leads to memory reclamation, which might degrade performance. If the active memory size is the same as the granted memory size, demand for memory is greater than the memory resources available. If the active memory is consistently low, the memory size might be too large.

If the host has enough free memory, check the resource shares, reservation, and limit of the virtual machines and resource pools on the host. Verify that the host settings are adequate and not lower than those set for the virtual machine.

If little free memory is available, or if you notice degradation in performance, consider taking the following actions.

**Table 1-86. Memory Performance Enhancement Advice**

#	Resolution
1	Verify that VMware Tools is installed on each virtual machine. The balloon driver is installed with VMware Tools and is critical to performance.
2	Verify that the balloon driver is enabled. The VMkernel regularly reclaims unused virtual machine memory by ballooning and swapping. Generally, this does not impact virtual machine performance.

**Table 1-86. Memory Performance Enhancement Advice (Continued)**

#	Resolution
3	Reduce the memory space on the virtual machine, and correct the cache size if it is too large. This frees up memory for other virtual machines.
4	If the memory reservation of the virtual machine is set to a value much higher than its active memory, decrease the reservation setting so that the VMkernel can reclaim the idle memory for other virtual machines on the host.
5	Migrate one or more virtual machines to a host in a DRS cluster.
6	Add physical memory to the host.

## Memory (MB)

The Memory (MB) chart displays memory data counters for resource pools or vApps.

### Description

This chart is located in the **Home** view of the resource pool or vApp **Performance** tab.

**Note** These data counter definitions are for virtual machines. At the resource pool level, the values are collected and totaled. The counter values in the chart represent the aggregate amounts of the virtual machine data. The counters that appear in the chart depend on the collection level set for your vCenter Server.

**Table 1-87. Data Counters**

Chart Label	Description
Active	<p>Sum of the active guest physical memory of all powered on virtual machines in the resource pool.</p> <ul style="list-style-type: none"> <li>▪ Counter: active</li> <li>▪ Stats Type: Absolute</li> <li>▪ Unit: Megabytes (MB)</li> <li>▪ Rollup Type: Average (Minimum/Maximum)</li> <li>▪ Collection Level: 2 (4)</li> </ul>
Balloon	<p>Sum of the guest physical memory reclaimed by the balloon driver for all powered on virtual machines in the resource pool.</p> <ul style="list-style-type: none"> <li>▪ Counter: vmmemctl</li> <li>▪ Stats Type: Absolute</li> <li>▪ Unit: Megabytes (MB)</li> <li>▪ Rollup Type: Average (Minimum/Maximum)</li> <li>▪ Collection Level: 1 (4)</li> </ul>

**Table 1-87. Data Counters (Continued)**

Chart Label	Description
Consumed	<p>Amount of the physical memory consumed by the virtual machine for the guest memory. Consumed memory does not include the overhead memory. It includes the shared memory and memory that might be reserved, but not actually used.</p> <p>consumed memory = memory granted – memory saved due to memory sharing</p> <ul style="list-style-type: none"> <li>■ Counter: consumed</li> <li>■ Stats Type: Absolute</li> <li>■ Unit: Megabytes (MB)</li> <li>■ Rollup Type: Average (Minimum/Maximum)</li> <li>■ Collection Level: 1 (4)</li> </ul>
Granted	<p>Sum of the guest physical memory granted for all powered on virtual machines. Granted memory is mapped to the host's machine memory.</p> <ul style="list-style-type: none"> <li>■ Counter: granted</li> <li>■ Stats Type: Absolute</li> <li>■ Unit: Megabytes (MB)</li> <li>■ Rollup Type: Average (Minimum/Maximum)</li> <li>■ Collection Level: 2 (4)</li> </ul>
Shared	<p>Amount of the guest physical memory shared with other virtual machines in the resource pool.</p>
Swapped	<p>Sum of the memory swapped by all powered on virtual machines in the resource pool.</p> <ul style="list-style-type: none"> <li>■ Counter: swapped</li> <li>■ Stats Type: Absolute</li> <li>■ Unit: Megabytes (MB)</li> <li>■ Rollup Type: Average (Minimum/Maximum)</li> <li>■ Collection Level: 2 (4)</li> </ul>

### Chart Analysis

A virtual machine's memory size must be slightly larger than the average guest memory usage. This enables the host to accommodate workload spikes without swapping memory among guests. Increasing the virtual machine memory size results in more overhead memory usage.

If sufficient swap space is available, a high balloon value does not cause performance problems. However, if the swapin and swapout values for the host are large, the host is probably lacking the amount of memory required to meet the demand.

If a virtual machine has high ballooning or swapping, check the amount of free physical memory on the host. A free memory value of 6% or less indicates that the host cannot meet the memory requirements. This leads to memory reclamation, which might degrade performance. If the active memory size is the same as the granted memory size, demand for memory is greater than the memory resources available. If the active memory is consistently low, the memory size might be too large.

If the host has enough free memory, check the resource shares, reservation, and limit of the virtual machines and resource pools on the host. Verify that the host settings are adequate and not lower than those set for the virtual machine.

If little free memory is available, or if you notice degradation in performance, consider taking the following actions.

**Table 1-88. Memory Performance Enhancement Advice**

#	Resolution
1	Verify that VMware Tools is installed on each virtual machine. The balloon driver is installed with VMware Tools and is critical to performance.
2	Verify that the balloon driver is enabled. The VMkernel regularly reclaims unused virtual machine memory by ballooning and swapping. Generally, this does not impact virtual machine performance.
3	Reduce the memory space on the virtual machine, and correct the cache size if it is too large. This frees up memory for other virtual machines.
4	If the memory reservation of the virtual machine is set to a value much higher than its active memory, decrease the reservation setting so that the VMkernel can reclaim the idle memory for other virtual machines on the host.
5	Migrate one or more virtual machines to a host in a DRS cluster.
6	Add physical memory to the host.

## vApps

The vApp charts contain information about CPU and memory usage for vApps. The help topic for each chart contains information about the data counters displayed in that chart. The counters available are determined by the collection level set for vCenter Server.

### CPU (MHz)

The CPU (MHz) chart displays CPU usage in the vApp or resource pool.

This chart is located in the **Home** view of the vApp or resource pool **Performance** tab.

### Counters

**Table 1-89. Data Counters**

Chart Label	Description
Usage	<p>CPU usage is the sum of the average CPU usage values of the virtual machines in the resource pool or vApp.</p> <p>CPU usage = number of cores * CPU frequency</p> <ul style="list-style-type: none"> <li>■ Counter: usagemhz</li> <li>■ Stats Type: Rate</li> <li>■ Unit: Megahertz (MHz)</li> <li>■ Rollup Type: Average (Minimum/Maximum)</li> <li>■ Collection Level: 1 (4)</li> </ul>



## Chart Analysis

A short spike in CPU usage indicates that you are making the best use of the resources available. However, if the value is constantly high, the CPU demanded is likely greater than the CPU capacity available. A high CPU usage value can lead to increased ready time and processor queuing of the virtual machines in the resource pool. Generally, if the CPU usage value for a virtual machine is above 90% and the CPU ready value for a virtual machine is above 20%, performance is impacted.

If performance is impacted, consider taking the following actions.

**Table 1-90. CPU Performance Enhancement Advice**

#	Resolution
1	Verify that VMware Tools is installed on each virtual machine.
2	Deploy single-threaded applications on uniprocessor virtual machines instead of SMP virtual machines.
3	Migrate one or more virtual machines to a new host.
4	Upgrade the physical CPUs or cores on each host if necessary.
5	Enable CPU-saving features, such as TCP Segmentation Offload.
6	Replace software I/O with dedicated hardware, such as iSCSI HBAs or TCP Segmentation Offload NICs.

## CPU Usage

The CPU Usage chart displays CPU usage of each virtual machine in the vApp or resource pool.

This chart is located in the **Virtual Machines** view of the vApp or resource pool **Performance** tab.

**Table 1-91. Data Counters**

Chart Label	Description
<i>virtual_machine</i>	Amount of CPU actively used by virtual machines. <ul style="list-style-type: none"> <li>■ Counter: usagemhz</li> <li>■ Stats Type: Rate</li> <li>■ Unit: Megahertz (MHz)</li> <li>■ Rollup Type: Average (Minimum/Maximum)</li> <li>■ Collection Level: 1 (4)</li> </ul>

## Chart Analysis

A short spike in CPU usage or CPU ready indicates that you are making the best use of the virtual machine resources. However, if the CPU usage value for a virtual machine is above 90% and the CPU ready value is above 20%, performance is being impacted.

If performance is impacted, consider taking the following actions.

**Table 1-92. CPU Performance Enhancement Advice**

#	Resolution
1	Verify that VMware Tools is installed on every virtual machine on the host.
2	Set the CPU reservations for all high-priority virtual machines to guarantee that they receive the CPU cycles required.

**Table 1-92. CPU Performance Enhancement Advice (Continued)**

#	Resolution
3	Compare the CPU usage value of a virtual machine with the CPU usage of other virtual machines on the host or in the resource pool. The stacked line chart on the host's <b>Virtual Machine</b> view shows the CPU usage for virtual machines on the host.
4	Determine whether the high ready time for the virtual machine resulted from its CPU usage time reaching the CPU limit setting. If so, increase the CPU limit on the virtual machine.
5	Increase the CPU shares to give the virtual machine more opportunities to run. The total ready time on the host might remain at the same level if the host system is constrained by CPU. If the host ready time doesn't decrease, set the CPU reservations for high-priority virtual machines to guarantee that they receive the required CPU cycles.
6	Increase the amount of memory allocated to the virtual machine. This decreases disk and or network activity for applications that cache. This might lower disk I/O and reduce the need for the host to virtualize the hardware. Virtual machines with smaller resource allocations generally accumulate more CPU ready time.
7	Reduce the number of virtual CPUs on a virtual machine to only the number required to execute the workload. For example, a single-threaded application on a four-way virtual machine only benefits from a single vCPU. But the hypervisor's maintenance of the three idle vCPUs takes CPU cycles that could be used for other work.
8	If the host is not already in a DRS cluster, add it to one. If the host is in a DRS cluster, increase the number of hosts and migrate one or more virtual machines onto the new host.
9	Upgrade the physical CPUs or cores on the host if necessary.
1	Use the newest version of hypervisor software, and enable CPU-saving features such as TCP Segmentation Offload, large
0	memory pages, and jumbo frames.

## Memory (MB)

The Memory (MB) chart displays memory usage in the vApp or resource pool.

This chart is located in the **Home** view of the vApp or resource pool **Performance** tab.

**Table 1-93. Data Counters**

Chart Label	Description
<i>resource_pool</i> or <i>vApp</i>	<p>Sum of the active memory used by all virtual machines in the resource pool or vApp. Active memory is determined by the VMkernel and includes overhead memory.</p> <p>memory usage = active memory / configured virtual machine memory size</p> <ul style="list-style-type: none"> <li>■ Counter: used</li> <li>■ Stats Type: Absolute</li> <li>■ Unit: Megabytes (MB)</li> <li>■ Rollup Type: Average</li> <li>■ Collection Level: 1</li> </ul>

## Chart Analysis

Memory usage is not an indicator of performance problems. Memory can be high if a host is swapping or ballooning, which can result in virtual machine guest swapping. In such cases, check for other problems, such as CPU over-commitment or storage latencies.

If you have constantly high memory usage in a cluster, resource pool, or vApp, consider taking the following actions.

**Table 1-94. Memory Performance Enhancement Advice**

#	Resolution
1	Verify that VMware Tools is installed on each virtual machine.
2	Verify that the balloon driver is enabled. The balloon driver is installed with VMware Tools and is critical to performance. The VMkernel regularly reclaims unused virtual machine memory by ballooning and swapping. Generally, it does not impact virtual machine performance.
3	If the balloon value is high, check the resource shares, reservations, and limits for the virtual machines and resource pools on the hosts. Verify that the host's settings are adequate and not lower than those set for the virtual machine. If free memory is available on the hosts and the virtual machines are experiencing high swap or balloon memory, the virtual machine (or resource pool, if it belongs to one) has reached its resource limit. Check the maximum resource limit set on that host.
4	If the cluster is not a DRS cluster, enable DRS. To enable DRS, perform the following tasks: <ol style="list-style-type: none"> <li>1 Select the cluster, and click the <b>Configure</b> tab.</li> <li>2 Under <b>Services</b>, click <b>vSphere DRS</b>.</li> <li>3 click <b>Edit</b>.  An Edit Cluster Settings dialog box opens.</li> <li>4 Click <b>Turn ON vSphere DRS</b>, and click <b>OK</b>.</li> </ol>
5	If the cluster is a DRS cluster: <ul style="list-style-type: none"> <li>■ Increase the number of hosts, and migrate one or more virtual machines to the new host.</li> <li>■ Check the aggressiveness threshold. If the value is low, increase the threshold. It might help avoid hot spots in the cluster.</li> </ul>
6	Add more physical memory to one or more hosts.

## Memory Consumed

The Memory Consumed chart displays the memory performance of the top ten virtual machines in the vApp or resource pool.

This chart is located in the **Virtual Machines** view of the vApp or resource pool **Performance** tab.

For resource pools and virtual machines in a resource pool or vApp, this chart is located in the **Resource Pools & Virtual Machines** view of the resource pool or vApp **Performance** tab.

**Table 1-95. Data Counters**

Chart Label	Description
<i>virtual_machine</i>	<p>Amount of host memory used by the virtual machine for its guest operating system's physical memory. Memory overhead is not included in consumed memory.</p> <p>consumed memory = memory granted - memory saved from page sharing</p> <p>For example, if a virtual machine has 100 MB of memory that is shared equally with three other virtual machines, its portion of the shared memory is 25 MB (100 MB ÷ 4 VMs). This amount is counted in the memory consumed data counter.</p> <ul style="list-style-type: none"> <li>■ Counter: consumed</li> <li>■ Stats Type: Absolute</li> <li>■ Unit: Megabytes (MB)</li> <li>■ Rollup Type: Average (Minimum/Maximum)</li> <li>■ Collection Level: 1 (4)</li> </ul>

## Chart Analysis

A virtual machine's memory size must be slightly larger than the average guest memory usage. This enables the host to accommodate workload spikes without swapping memory among guests. Increasing the virtual machine memory size results in more overhead memory usage.

If sufficient swap space is available, a high balloon value does not cause performance problems. However, if the swpin and swpout values for the host are large, the host is probably lacking the amount of memory required to meet the demand.

If a virtual machine has high ballooning or swapping, check the amount of free physical memory on the host. A free memory value of 6% or less indicates that the host cannot meet the memory requirements. This leads to memory reclamation, which might degrade performance. If the active memory size is the same as the granted memory size, demand for memory is greater than the memory resources available. If the active memory is consistently low, the memory size might be too large.

If the host has enough free memory, check the resource shares, reservation, and limit of the virtual machines and resource pools on the host. Verify that the host settings are adequate and not lower than those set for the virtual machine.

If little free memory is available, or if you notice degradation in performance, consider taking the following actions.

**Table 1-96. Memory Performance Enhancement Advice**

#	Resolution
1	Verify that VMware Tools is installed on each virtual machine. The balloon driver is installed with VMware Tools and is critical to performance.
2	Verify that the balloon driver is enabled. The VMkernel regularly reclaims unused virtual machine memory by ballooning and swapping. Generally, this does not impact virtual machine performance.
3	Reduce the memory space on the virtual machine, and correct the cache size if it is too large. This frees up memory for other virtual machines.
4	If the memory reservation of the virtual machine is set to a value much higher than its active memory, decrease the reservation setting so that the VMkernel can reclaim the idle memory for other virtual machines on the host.
5	Migrate one or more virtual machines to a host in a DRS cluster.
6	Add physical memory to the host.

## Virtual Machines

The virtual machine charts contain information about CPU, disk, memory, network, storage, and fault tolerance for virtual machines. The help topic for each chart contains information about the data counters displayed in that chart. The counters available are determined by the collection level set for vCenter Server.

### CPU (%)

The CPU (%) chart displays virtual machine CPU usage and ready values.

This chart is located in the **Home** view of the virtual machine **Performance** tab.

**Table 1-97. Data Counters**

Chart Label	Description
Usage	<p>Amount of actively used virtual CPU as a percentage of total available CPU.</p> <p>CPU usage is the average CPU utilization over all available virtual CPUs in the virtual machine.</p> <p>For example, if a virtual machine with one virtual CPU is running on a host that has four physical CPUs and the CPU usage is 100%, the virtual machine is using one physical CPU completely.</p> <p>virtual CPU usage = <math>\text{usagemhz} \div (\text{number of virtual CPUs} \times \text{core frequency})</math></p> <hr/> <p><b>Note</b> This is the host's view of the CPU usage, not the guest operating system view.</p> <hr/> <ul style="list-style-type: none"> <li>■ Counter: usage</li> <li>■ Stats Type: Rate</li> <li>■ Unit: Percentage (%). Precision is to 1/100%. A value between 0 and 100.</li> <li>■ Rollup Type: Average (Minimum/Maximum)</li> <li>■ Collection Level: 1 (4)</li> </ul>
Ready	<p>Percentage of time that the virtual machine was ready, but could not get scheduled to run on the physical CPU.</p> <p>CPU ready time is dependent on the number of virtual machines on the host and their CPU loads. At collection level 1, the average CPU ready time of all virtual CPUs on the virtual machine is displayed. At collection level 3, the average CPU ready time of each virtual CPU is also displayed.</p> <ul style="list-style-type: none"> <li>■ Counter: ready</li> <li>■ Stats Type: Rate</li> <li>■ Unit: Percentage (%)</li> <li>■ Rollup Type: Summation</li> <li>■ Collection Level: 1</li> </ul>

### Chart Analysis

A short spike in CPU usage or CPU ready indicates that you are making the best use of the virtual machine resources. However, if the CPU usage value for a virtual machine is above 90% and the CPU ready value is above 20%, performance is being impacted.

If performance is impacted, consider taking the following actions.

**Table 1-98. CPU Performance Enhancement Advice**

#	Resolution
1	Verify that VMware Tools is installed on every virtual machine on the host.
2	Set the CPU reservations for all high-priority virtual machines to guarantee that they receive the CPU cycles required.
3	Compare the CPU usage value of a virtual machine with the CPU usage of other virtual machines on the host or in the resource pool. The stacked line chart on the host's <b>Virtual Machine</b> view shows the CPU usage for virtual machines on the host.
4	Determine whether the high ready time for the virtual machine resulted from its CPU usage time reaching the CPU limit setting. If so, increase the CPU limit on the virtual machine.
5	Increase the CPU shares to give the virtual machine more opportunities to run. The total ready time on the host might remain at the same level if the host system is constrained by CPU. If the host ready time doesn't decrease, set the CPU reservations for high-priority virtual machines to guarantee that they receive the required CPU cycles.

**Table 1-98. CPU Performance Enhancement Advice (Continued)**

#	Resolution
6	Increase the amount of memory allocated to the virtual machine. This decreases disk and or network activity for applications that cache. This might lower disk I/O and reduce the need for the host to virtualize the hardware. Virtual machines with smaller resource allocations generally accumulate more CPU ready time.
7	Reduce the number of virtual CPUs on a virtual machine to only the number required to execute the workload. For example, a single-threaded application on a four-way virtual machine only benefits from a single vCPU. But the hypervisor's maintenance of the three idle vCPUs takes CPU cycles that could be used for other work.
8	If the host is not already in a DRS cluster, add it to one. If the host is in a DRS cluster, increase the number of hosts and migrate one or more virtual machines onto the new host.
9	Upgrade the physical CPUs or cores on the host if necessary.
10	Use the newest version of hypervisor software, and enable CPU-saving features such as TCP Segmentation Offload, large memory pages, and jumbo frames.

## CPU Usage (MHz)

The CPU Usage (MHz) chart displays virtual machine CPU usage.

This chart is located in the **Home** view of the Virtual Machine **Performance** tab.

**Table 1-99. Data Counters**

Chart Label	Description
Usage	Amount of actively used virtual CPU.
	<b>Note</b> The host's view of the CPU usage, not the guest operating system view.
	<ul style="list-style-type: none"> <li>■ Counter: usagemhz</li> <li>■ Stats Type: rate</li> <li>■ Unit: MegaHertz (MHz)</li> <li>■ Rollup Type: Average (Minimum/Maximum)</li> <li>■ Collection Level: 1 (4)</li> </ul>

## Chart Analysis

A short spike in CPU usage or CPU ready indicates that you are making the best use of the virtual machine resources. However, if the CPU usage value for a virtual machine is above 90% and the CPU ready value is above 20%, performance is being impacted.

If performance is impacted, consider taking the following actions.

**Table 1-100. CPU Performance Enhancement Advice**

#	Resolution
1	Verify that VMware Tools is installed on every virtual machine on the host.
2	Set the CPU reservations for all high-priority virtual machines to guarantee that they receive the CPU cycles required.
3	Compare the CPU usage value of a virtual machine with the CPU usage of other virtual machines on the host or in the resource pool. The stacked line chart on the host's <b>Virtual Machine</b> view shows the CPU usage for virtual machines on the host.

**Table 1-100. CPU Performance Enhancement Advice (Continued)**

#	Resolution
4	Determine whether the high ready time for the virtual machine resulted from its CPU usage time reaching the CPU limit setting. If so, increase the CPU limit on the virtual machine.
5	Increase the CPU shares to give the virtual machine more opportunities to run. The total ready time on the host might remain at the same level if the host system is constrained by CPU. If the host ready time doesn't decrease, set the CPU reservations for high-priority virtual machines to guarantee that they receive the required CPU cycles.
6	Increase the amount of memory allocated to the virtual machine. This decreases disk and or network activity for applications that cache. This might lower disk I/O and reduce the need for the host to virtualize the hardware. Virtual machines with smaller resource allocations generally accumulate more CPU ready time.
7	Reduce the number of virtual CPUs on a virtual machine to only the number required to execute the workload. For example, a single-threaded application on a four-way virtual machine only benefits from a single vCPU. But the hypervisor's maintenance of the three idle vCPUs takes CPU cycles that could be used for other work.
8	If the host is not already in a DRS cluster, add it to one. If the host is in a DRS cluster, increase the number of hosts and migrate one or more virtual machines onto the new host.
9	Upgrade the physical CPUs or cores on the host if necessary.
10	Use the newest version of hypervisor software, and enable CPU-saving features such as TCP Segmentation Offload, large memory pages, and jumbo frames.

## Disk (KBps)

The Disk (KBps) chart displays disk usage for the virtual machine.

It is located in the **Home** view of the virtual machine **Performance** tab.

**Table 1-101. Data Counters**

Chart Label	Description
Usage	<p>Average data I/O rate across all virtual disks on the virtual machine.</p> <ul style="list-style-type: none"> <li>■ Counter: usage</li> <li>■ Stats Type: Rate</li> <li>■ Unit: Kilobytes per second (KBps)</li> <li>■ Rollup Type: Average (Minimum/Maximum)</li> <li>■ Collection Level: 1 (4)</li> </ul>

## Chart Analysis

Use the disk charts to monitor average disk loads and to determine trends in disk usage. For example, you might notice a performance degradation with applications that frequently read from and write to the hard disk. If you see a spike in the number of disk read or write requests, check whether any such applications were running then.

The best ways to determine if your vSphere environment is experiencing disk problems is to monitor the disk latency data counters. You can use the advanced performance charts to view these statistics.

- The `kernelLatency` data counter measures the average amount of time, in milliseconds, that the VMkernel spends processing each SCSI command. For best performance, the value must be 0 -1 milliseconds. If the value is greater than 4 ms, the virtual machines on the host are trying to send more throughput to the storage system than the configuration supports. Check the CPU usage, and increase the queue depth.
- The `deviceLatency` data counter measures the average amount of time, in milliseconds, to complete a SCSI command from the physical device. Depending on your hardware, a number greater than 15 ms indicates probable problems with the storage array. Move the active VMDK to a volume with more spindles or add disks to the LUN.
- The `queueLatency` data counter measures the average amount of time taken per SCSI command in the VMkernel queue. This value must always be zero. If not, the workload is too high and the array cannot process the data fast enough.

If the disk latency values are high, or if you notice other problems with disk I/O performance, consider taking the following actions.

**Table 1-102. Disk I/O Performance Enhancement Advice**

#	Resolution
1	<p>Increase the virtual machine memory. It allows more operating system caching, which reduces I/O activity. Note: It might require you to increase the host memory. Increasing memory might reduce the need to store data because databases can utilize the system memory to cache data and avoid disk access.</p> <p>To verify that virtual machines have adequate memory, check swap statistics in the guest operating system. Increase the guest memory, but not to an extent that leads to excessive host memory swapping. Install VMware Tools so that memory ballooning can occur.</p>
2	Defragment the file systems on all guests.
3	Disable antivirus on-demand scans on the VMDK and VMEM files.
4	Use the vendor's array tools to determine the array performance statistics. When too many servers simultaneously access common elements on an array, the disks might have trouble keeping up. To increase throughput, consider array-side improvements.
5	Use Storage vMotion to migrate I/O-intensive virtual machines across multiple hosts.
6	Balance the disk load across all physical resources available. Spread heavily used storage across LUNs that are accessed by different adapters. Use separate queues for each adapter to improve disk efficiency.
7	Configure the HBAs and RAID controllers for optimal use. Verify that the queue depths and cache settings on the RAID controllers are adequate. If not, increase the number of outstanding disk requests for the virtual machine by adjusting the <code>Disk.SchedNumReqOutstanding</code> parameter. For more information, see <i>vSphere Storage</i> .
8	For resource-intensive virtual machines, separate the virtual machine's physical disk drive from the drive with the system page file. It alleviates disk spindle contention during periods of high use.
9	On systems with sizable RAM, disable memory trimming by adding the line <code>MemTrimRate=0</code> to the virtual machine's VMX file.
10	If the combined disk I/O is higher than a single HBA capacity, use multipathing or multiple links.



**Table 1-102. Disk I/O Performance Enhancement Advice (Continued)**

#	Resolution
11	For ESXi hosts, create virtual disks as preallocated. When you create a virtual disk for a guest operating system, select <b>Allocate all disk space now</b> . The performance degradation associated with reassigning additional disk space does not occur, and the disk is less likely to become fragmented.
12	Use the most current host hardware.

## Disk Rate (KBps)

The Disk Rate chart displays disk usage for the virtual machine.

This chart is located in the **Home** view of the virtual machine **Performance** tab. It is available only at collection levels 3 and 4.

**Table 1-103. Data Counters**

Chart Label	Description
Read	<p>Number of disk read commands completed on each virtual disk on the virtual machine, per second. The aggregate number of all disk read commands per second is also displayed in the chart.</p> <p>Read rate = blocksRead per second × blockSize</p> <ul style="list-style-type: none"> <li>■ Counter: read</li> <li>■ Stats Type: Rate</li> <li>■ Unit: Kilobytes per second (KBps)</li> <li>■ Rollup Type: Average</li> <li>■ Collection Level: 2</li> </ul>
Write	<p>Number of disk write commands completed on each virtual disk on the virtual machine, per second. The aggregate number of all disk write commands per second is also displayed in the chart.</p> <p>Write rate = blocksWritten per second × blockSize</p> <ul style="list-style-type: none"> <li>■ Counter: write</li> <li>■ Stats Type: Rate</li> <li>■ Unit: Kilobytes per second (KBps)</li> <li>■ Rollup Type: Average</li> <li>■ Collection Level: 2</li> </ul>

## Chart Analysis

Use the disk charts to monitor average disk loads and to determine trends in disk usage. For example, you might notice a performance degradation with applications that frequently read from and write to the hard disk. If you see a spike in the number of disk read or write requests, check whether any such applications were running then.

The best ways to determine if your vSphere environment is experiencing disk problems is to monitor the disk latency data counters. You can use the advanced performance charts to view these statistics.

- The `kernelLatency` data counter measures the average amount of time, in milliseconds, that the VMkernel spends processing each SCSI command. For best performance, the value must be 0 -1 milliseconds. If the value is greater than 4 ms, the virtual machines on the host are trying to send more throughput to the storage system than the configuration supports. Check the CPU usage, and increase the queue depth.
- The `deviceLatency` data counter measures the average amount of time, in milliseconds, to complete a SCSI command from the physical device. Depending on your hardware, a number greater than 15 ms indicates probable problems with the storage array. Move the active VMDK to a volume with more spindles or add disks to the LUN.
- The `queueLatency` data counter measures the average amount of time taken per SCSI command in the VMkernel queue. This value must always be zero. If not, the workload is too high and the array cannot process the data fast enough.

If the disk latency values are high, or if you notice other problems with disk I/O performance, consider taking the following actions.

**Table 1-104. Disk I/O Performance Enhancement Advice**

#	Resolution
1	<p>Increase the virtual machine memory. It allows more operating system caching, which reduces I/O activity. Note: It might require you to increase the host memory. Increasing memory might reduce the need to store data because databases can utilize the system memory to cache data and avoid disk access.</p> <p>To verify that virtual machines have adequate memory, check swap statistics in the guest operating system. Increase the guest memory, but not to an extent that leads to excessive host memory swapping. Install VMware Tools so that memory ballooning can occur.</p>
2	Defragment the file systems on all guests.
3	Disable antivirus on-demand scans on the VMDK and VMEM files.
4	Use the vendor's array tools to determine the array performance statistics. When too many servers simultaneously access common elements on an array, the disks might have trouble keeping up. To increase throughput, consider array-side improvements.
5	Use Storage vMotion to migrate I/O-intensive virtual machines across multiple hosts.
6	Balance the disk load across all physical resources available. Spread heavily used storage across LUNs that are accessed by different adapters. Use separate queues for each adapter to improve disk efficiency.
7	Configure the HBAs and RAID controllers for optimal use. Verify that the queue depths and cache settings on the RAID controllers are adequate. If not, increase the number of outstanding disk requests for the virtual machine by adjusting the <code>Disk.SchedNumReqOutstanding</code> parameter. For more information, see <i>vSphere Storage</i> .
8	For resource-intensive virtual machines, separate the virtual machine's physical disk drive from the drive with the system page file. It alleviates disk spindle contention during periods of high use.
9	On systems with sizable RAM, disable memory trimming by adding the line <code>MemTrimRate=0</code> to the virtual machine's VMX file.
10	If the combined disk I/O is higher than a single HBA capacity, use multipathing or multiple links.

**Table 1-104. Disk I/O Performance Enhancement Advice (Continued)**

#	Resolution
11	For ESXi hosts, create virtual disks as preallocated. When you create a virtual disk for a guest operating system, select <b>Allocate all disk space now</b> . The performance degradation associated with reassigning additional disk space does not occur, and the disk is less likely to become fragmented.
12	Use the most current host hardware.

## Disk Requests (Number)

The Disk Requests chart displays disk usage for the virtual machine.

This chart is located in the **Home** view of the virtual machine **Performance** tab. It is available only at collection levels 3 and 4.

**Table 1-105. Data Counters**

Chart Label	Description
Read Requests	<p>Number of disk read commands completed on each virtual disk on the virtual machine. The aggregate number of all disk read commands is also displayed in the chart.</p> <ul style="list-style-type: none"> <li>■ Counter: numberRead</li> <li>■ Stats Type: Absolute</li> <li>■ Unit: Number</li> <li>■ Rollup Type: Summation</li> <li>■ Collection Level: 3</li> </ul>
Write Requests	<p>Number of disk write commands completed on each virtual disk on the virtual machine. The aggregate number of all disk write commands is also displayed in the chart.</p> <ul style="list-style-type: none"> <li>■ Counter: numberWrite</li> <li>■ Stats Type: Absolute</li> <li>■ Unit: Number</li> <li>■ Rollup Type: Summation</li> <li>■ Collection Level: 3</li> </ul>

## Chart Analysis

Use the disk charts to monitor average disk loads and to determine trends in disk usage. For example, you might notice a performance degradation with applications that frequently read from and write to the hard disk. If you see a spike in the number of disk read or write requests, check whether any such applications were running then.

The best ways to determine if your vSphere environment is experiencing disk problems is to monitor the disk latency data counters. You can use the advanced performance charts to view these statistics.

- The `kernelLatency` data counter measures the average amount of time, in milliseconds, that the VMkernel spends processing each SCSI command. For best performance, the value must be 0 -1 milliseconds. If the value is greater than 4 ms, the virtual machines on the host are trying to send more throughput to the storage system than the configuration supports. Check the CPU usage, and increase the queue depth.
- The `deviceLatency` data counter measures the average amount of time, in milliseconds, to complete a SCSI command from the physical device. Depending on your hardware, a number greater than 15 ms indicates probable problems with the storage array. Move the active VMDK to a volume with more spindles or add disks to the LUN.
- The `queueLatency` data counter measures the average amount of time taken per SCSI command in the VMkernel queue. This value must always be zero. If not, the workload is too high and the array cannot process the data fast enough.

If the disk latency values are high, or if you notice other problems with disk I/O performance, consider taking the following actions.

**Table 1-106. Disk I/O Performance Enhancement Advice**

#	Resolution
1	<p>Increase the virtual machine memory. It allows more operating system caching, which reduces I/O activity. Note: It might require you to increase the host memory. Increasing memory might reduce the need to store data because databases can utilize the system memory to cache data and avoid disk access.</p> <p>To verify that virtual machines have adequate memory, check swap statistics in the guest operating system. Increase the guest memory, but not to an extent that leads to excessive host memory swapping. Install VMware Tools so that memory ballooning can occur.</p>
2	Defragment the file systems on all guests.
3	Disable antivirus on-demand scans on the VMDK and VMEM files.
4	Use the vendor's array tools to determine the array performance statistics. When too many servers simultaneously access common elements on an array, the disks might have trouble keeping up. To increase throughput, consider array-side improvements.
5	Use Storage vMotion to migrate I/O-intensive virtual machines across multiple hosts.
6	Balance the disk load across all physical resources available. Spread heavily used storage across LUNs that are accessed by different adapters. Use separate queues for each adapter to improve disk efficiency.
7	Configure the HBAs and RAID controllers for optimal use. Verify that the queue depths and cache settings on the RAID controllers are adequate. If not, increase the number of outstanding disk requests for the virtual machine by adjusting the <code>Disk.SchedNumReqOutstanding</code> parameter. For more information, see <i>vSphere Storage</i> .
8	For resource-intensive virtual machines, separate the virtual machine's physical disk drive from the drive with the system page file. It alleviates disk spindle contention during periods of high use.
9	On systems with sizable RAM, disable memory trimming by adding the line <code>MemTrimRate=0</code> to the virtual machine's VMX file.
10	If the combined disk I/O is higher than a single HBA capacity, use multipathing or multiple links.

**Table 1-106. Disk I/O Performance Enhancement Advice (Continued)**

#	Resolution
11	For ESXi hosts, create virtual disks as preallocated. When you create a virtual disk for a guest operating system, select <b>Allocate all disk space now</b> . The performance degradation associated with reassigning additional disk space does not occur, and the disk is less likely to become fragmented.
12	Use the most current host hardware.

## Virtual Disk Requests (Number)

The Virtual Disk Requests chart displays virtual disk usage for the virtual machine.

After you click **Overview** on the **Performance** tab of the virtual machine, you can view this chart by selecting **Home** from the **View** drop-down menu. It is available at collection (display) levels 3 and 4.

**Table 1-107. Data Counters**

Chart Label	Description
Read Requests	<p>Number of virtual disk read commands completed on each virtual disk on the virtual machine. The aggregate number of all virtual disk read commands is also displayed in the chart.</p> <ul style="list-style-type: none"> <li>▪ Counter: numberRead</li> <li>▪ Stats Type: Absolute</li> <li>▪ Unit: Number</li> <li>▪ Rollup Type: Average</li> <li>▪ Collection Level: 2</li> </ul>
Write Requests	<p>Number of virtual disk write commands completed on each virtual disk on the virtual machine. The aggregate number of all virtual disk write commands is also displayed in the chart.</p> <ul style="list-style-type: none"> <li>▪ Counter: numberWrite</li> <li>▪ Stats Type: Absolute</li> <li>▪ Unit: Number</li> <li>▪ Rollup Type: Average</li> <li>▪ Collection Level: 2</li> </ul>

## Virtual Disk Rate (KBps)

The Virtual Disk Rate chart displays virtual disk usage rate for the virtual machine.

After you click **Overview** on the **Performance** tab of the virtual machine, you can view this chart by selecting **Home** from the **View** drop-down menu. It is available only at collection levels 3 and 4.

**Table 1-108. Data Counters**

Chart Label	Description
Read Requests	<p>Number of virtual disk read commands completed on each virtual disk on the virtual machine. The aggregate number of all virtual disk read commands per second is also displayed in the chart. Read rate = blocksRead per second × blockSize</p> <ul style="list-style-type: none"> <li>■ Counter: read</li> <li>■ Stats Type: Rate</li> <li>■ Unit: KiloBytes per second (KBps)</li> <li>■ Rollup Type: Average</li> <li>■ Collection Level: 3</li> </ul>
Write Requests	<p>Number of virtual disk write commands completed on each virtual disk on the virtual machine per second. The aggregate number of all virtual disk write commands per second is also displayed in the chart. Write rate = blocksWritten per second × blockSize</p> <ul style="list-style-type: none"> <li>■ Counter: write</li> <li>■ Stats Type: Rate</li> <li>■ Unit: KiloBytes per second (KBps)</li> <li>■ Rollup Type: Average</li> <li>■ Collection Level: 3</li> </ul>

## Memory (%)

The Memory (%) chart monitors virtual machine memory usage.

This chart is located in the **Home** view of the virtual machine **Performance** tab.

## Virtual Machine Counters

**Note** Guest physical memory refers to the virtual hardware memory presented to a virtual machine for its guest operating system.

**Table 1-109. Data Counters**

Chart Label	Description
Usage	<p>Amount of guest physical memory currently in use on the virtual machine.</p> <ul style="list-style-type: none"> <li>■ Counter: usage</li> <li>■ Stats Type: Absolute</li> <li>■ Unit: Percentage (%)</li> <li>■ Rollup Type: Average (Minimum/Maximum)</li> <li>■ Collection Level: 1 (4)</li> </ul>

## Chart Analysis

A virtual machine's memory size must be slightly larger than the average guest memory usage. This enables the host to accommodate workload spikes without swapping memory among guests. Increasing the virtual machine memory size results in more overhead memory usage.

If sufficient swap space is available, a high balloon value does not cause performance problems. However, if the `swpin` and `swpout` values for the host are large, the host is probably lacking the amount of memory required to meet the demand.

If a virtual machine has high ballooning or swapping, check the amount of free physical memory on the host. A free memory value of 6% or less indicates that the host cannot meet the memory requirements. This leads to memory reclamation, which might degrade performance. If the active memory size is the same as the granted memory size, demand for memory is greater than the memory resources available. If the active memory is consistently low, the memory size might be too large.

If the host has enough free memory, check the resource shares, reservation, and limit of the virtual machines and resource pools on the host. Verify that the host settings are adequate and not lower than those set for the virtual machine.

If little free memory is available, or if you notice degradation in performance, consider taking the following actions.

**Table 1-110. Memory Performance Enhancement Advice**

#	Resolution
1	Verify that VMware Tools is installed on each virtual machine. The balloon driver is installed with VMware Tools and is critical to performance.
2	Verify that the balloon driver is enabled. The VMkernel regularly reclaims unused virtual machine memory by ballooning and swapping. Generally, this does not impact virtual machine performance.
3	Reduce the memory space on the virtual machine, and correct the cache size if it is too large. This frees up memory for other virtual machines.
4	If the memory reservation of the virtual machine is set to a value much higher than its active memory, decrease the reservation setting so that the VMkernel can reclaim the idle memory for other virtual machines on the host.
5	Migrate one or more virtual machines to a host in a DRS cluster.
6	Add physical memory to the host.

## Memory (MB)

The Memory (MB) chart displays virtual machine balloon memory.

This chart is located in the **Home** view of the virtual machine **Performance** tab.

**Table 1-111. Data Counters**

Chart Label	Description
Balloon	Amount of guest physical memory reclaimed from the virtual machine by the balloon driver. <ul style="list-style-type: none"> <li>■ Counter: <code>vmmemctl</code></li> <li>■ Stats Type: Absolute</li> <li>■ Unit: Megabytes (MB)</li> <li>■ Rollup Type: Average (Minimum/Maximum)</li> <li>■ Collection Level: 1 (4)</li> </ul>

## Chart Analysis

A virtual machine's memory size must be slightly larger than the average guest memory usage. This enables the host to accommodate workload spikes without swapping memory among guests. Increasing the virtual machine memory size results in more overhead memory usage.

If sufficient swap space is available, a high balloon value does not cause performance problems. However, if the `swpin` and `swpout` values for the host are large, the host is probably lacking the amount of memory required to meet the demand.

If a virtual machine has high ballooning or swapping, check the amount of free physical memory on the host. A free memory value of 6% or less indicates that the host cannot meet the memory requirements. This leads to memory reclamation, which might degrade performance. If the active memory size is the same as the granted memory size, demand for memory is greater than the memory resources available. If the active memory is consistently low, the memory size might be too large.

If the host has enough free memory, check the resource shares, reservation, and limit of the virtual machines and resource pools on the host. Verify that the host settings are adequate and not lower than those set for the virtual machine.

If little free memory is available, or if you notice degradation in performance, consider taking the following actions.

**Table 1-112. Memory Performance Enhancement Advice**

#	Resolution
1	Verify that VMware Tools is installed on each virtual machine. The balloon driver is installed with VMware Tools and is critical to performance.
2	Verify that the balloon driver is enabled. The VMkernel regularly reclaims unused virtual machine memory by ballooning and swapping. Generally, this does not impact virtual machine performance.
3	Reduce the memory space on the virtual machine, and correct the cache size if it is too large. This frees up memory for other virtual machines.
4	If the memory reservation of the virtual machine is set to a value much higher than its active memory, decrease the reservation setting so that the VMkernel can reclaim the idle memory for other virtual machines on the host.
5	Migrate one or more virtual machines to a host in a DRS cluster.
6	Add physical memory to the host.

## Memory (MBps)

The Memory (MBps) chart displays virtual machine memory swap rates.

This chart is located in the **Home** view of the Virtual Machine **Performance** tab.



**Table 1-113. Data Counters**

Chart Label	Description
swapinRate	<p>Average rate at which memory is swapped into the virtual machine.</p> <ul style="list-style-type: none"> <li>■ Counter: swapinRate</li> <li>■ Stats Type: Rate</li> <li>■ Unit: MegaBytes per second (MBps)</li> <li>■ Rollup Type: Average (Minimum/Maximum)</li> <li>■ Collection Level: 1 (4)</li> </ul>
swapoutRate	<p>Average rate at which memory is swapped out of the virtual machine.</p> <ul style="list-style-type: none"> <li>■ Counter: swapoutRate</li> <li>■ Stats Type: Rate</li> <li>■ Unit: MegaBytes per second (MBps)</li> <li>■ Rollup Type: Average (Minimum/Maximum)</li> <li>■ Collection Level: 1 (4)</li> </ul>

### Chart Analysis

A virtual machine's memory size must be slightly larger than the average guest memory usage. This enables the host to accommodate workload spikes without swapping memory among guests. Increasing the virtual machine memory size results in more overhead memory usage.

If there is sufficient swap space, a high balloon value is not a performance issue. However, if the swapin and swapout values for the host are large, the host is probably lacking the memory required to meet the demand.

If a virtual machine has high ballooning or swapping, check the amount of free physical memory on the host. The host might require more memory resources. If it does not, check the resource shares, reservation, and limit of the virtual machines and resource pools on the host. Verify that the host settings are adequate and not lower than those set for the virtual machine.

If memory usage is high or you notice degradation in performance, consider taking the following actions.

**Table 1-114. Memory Performance Enhancement Advice**

#	Resolution
1	Verify that VMware Tools is installed on each virtual machine. The balloon driver is installed with VMware Tools and is critical to performance.
2	Verify that the balloon driver is enabled. The VMkernel regularly reclaims unused virtual machine memory by ballooning and swapping. Generally, this does not impact virtual machine performance.
3	Reduce the memory space on the virtual machine, and correct the cache size if it is too large. This frees up memory for other virtual machines.
4	If the memory reservation of a virtual machine is set to a value much higher than its active memory, decrease the reservation setting so that the VMkernel can reclaim the idle memory for other virtual machines on the host.
5	Migrate one or more virtual machines to a host in a DRS cluster.
6	Add physical memory to the host.

## Memory (MB)

The Memory (MB) chart displays memory data counters for virtual machines.

This chart is located in the **Home** view of the virtual machine **Performance** tab. It appears only at collection levels 2, 3, and 4.

In the following descriptions, the guest physical memory refers to the virtual-hardware memory presented to a virtual machine for its guest operating system. Machine memory is actual physical RAM in the host. Note that not all counters are collected at collection level 1.

**Table 1-115. Data Counters**

Chart Label	Description
Active	<p>Amount of guest physical memory in use by the virtual machine.</p> <p>Active memory is estimated by VMkernel statistical sampling and represents the actual amount of memory the virtual machine needs. The value is based on the current workload of the virtual machine.</p> <ul style="list-style-type: none"> <li>■ Counter: active</li> <li>■ Stats Type: Absolute</li> <li>■ Unit: Megabytes (MB)</li> <li>■ Rollup Type: Average (Minimum/Maximum)</li> <li>■ Collection Level: 2 (4)</li> </ul>
Balloon	<p>Amount of guest physical memory reclaimed from the virtual machine by the balloon driver.</p> <ul style="list-style-type: none"> <li>■ Counter: vmmemctl</li> <li>■ Stats Type: Absolute</li> <li>■ Unit: Megabytes (MB)</li> <li>■ Rollup Type: Average (Minimum/Maximum)</li> <li>■ Collection Level: 1 (4)</li> </ul>
Balloon Target	<p>Desired amount of virtual machine balloon memory.</p> <p>Balloon target memory is estimated by the VMkernel.</p> <p>If the balloon target amount is greater than the balloon amount, the VMkernel inflates the balloon amount, which reclaims more virtual machine memory. If the balloon target amount is less than the balloon amount, the VMkernel deflates the balloon, which allows the virtual machine to reallocate memory when needed.</p> <ul style="list-style-type: none"> <li>■ Counter: vmmemctltarget</li> <li>■ Stats Type: Absolute</li> <li>■ Unit: Megabytes (MB)</li> <li>■ Rollup Type: Average (Minimum/Maximum)</li> <li>■ Collection Level: 2 (4)</li> </ul>

**Table 1-115. Data Counters (Continued)**

Chart Label	Description
Consumed	<p>Amount of guest physical memory consumed by the virtual machine for guest memory. Consumed memory does not include overhead memory. It includes shared memory and memory that might be reserved, but not actually used.</p> <p>consumed memory = memory granted - memory saved due to memory sharing</p> <ul style="list-style-type: none"> <li>■ Counter: consumed</li> <li>■ Stats Type: Absolute</li> <li>■ Unit: Megabytes (MB)</li> <li>■ Rollup Type: Average (Minimum/Maximum)</li> <li>■ Collection Level: 1 (4)</li> </ul>
Shared	<p>Amount of guest physical memory available for sharing. Memory sharing occurs through transparent page sharing.</p> <ul style="list-style-type: none"> <li>■ Counter: shared</li> <li>■ Stats Type: Absolute</li> <li>■ Unit: Megabytes (MB)</li> <li>■ Rollup Type: Average (Minimum/Maximum)</li> <li>■ Collection Level: 2 (4)</li> </ul>
Swapped	<p>The amount of guest physical memory swapped out to the disk by the VMkernel. This data counter measures VMkernel swapping and not to guest OS swapping.</p> <p>swapped = swapout – swapin</p> <p><b>Note</b> In some cases, vMotion can skew these values and cause a virtual machine to arrive on a host with some memory already swapped out. As a result, the swapped value can be greater than the swapout – swapin value.</p> <ul style="list-style-type: none"> <li>■ Counter: swapped</li> <li>■ Stats Type: Absolute</li> <li>■ Unit: Megabytes (MB)</li> <li>■ Rollup Type: Average (Minimum/Maximum)</li> <li>■ Collection Level: 2 (4)</li> </ul>

## Chart Analysis

A virtual machine's memory size must be slightly larger than the average guest memory usage. This enables the host to accommodate workload spikes without swapping memory among guests. Increasing the virtual machine memory size results in more overhead memory usage.

If sufficient swap space is available, a high balloon value does not cause performance problems. However, if the swapin and swapout values for the host are large, the host is probably lacking the amount of memory required to meet the demand.

If a virtual machine has high ballooning or swapping, check the amount of free physical memory on the host. A free memory value of 6% or less indicates that the host cannot meet the memory requirements. This leads to memory reclamation, which might degrade performance. If the active memory size is the same as the granted memory size, demand for memory is greater than the memory resources available. If the active memory is consistently low, the memory size might be too large.

If the host has enough free memory, check the resource shares, reservation, and limit of the virtual machines and resource pools on the host. Verify that the host settings are adequate and not lower than those set for the virtual machine.

If little free memory is available, or if you notice degradation in performance, consider taking the following actions.

**Table 1-116. Memory Performance Enhancement Advice**

#	Resolution
1	Verify that VMware Tools is installed on each virtual machine. The balloon driver is installed with VMware Tools and is critical to performance.
2	Verify that the balloon driver is enabled. The VMkernel regularly reclaims unused virtual machine memory by ballooning and swapping. Generally, this does not impact virtual machine performance.
3	Reduce the memory space on the virtual machine, and correct the cache size if it is too large. This frees up memory for other virtual machines.
4	If the memory reservation of the virtual machine is set to a value much higher than its active memory, decrease the reservation setting so that the VMkernel can reclaim the idle memory for other virtual machines on the host.
5	Migrate one or more virtual machines to a host in a DRS cluster.
6	Add physical memory to the host.

## Network (Mbps)

The Network (Mbps) chart displays network bandwidth for the virtual machine.

This chart is located in the **Home** view of the Virtual Machine **Performance** tab.

**Table 1-117. Virtual Machine Counters**

Chart Label	Description
Usage	<p>Average rate at which data is transmitted and received across all virtual NIC instances connected to the virtual machine.</p> <ul style="list-style-type: none"> <li>■ Counter: usage</li> <li>■ Stats Type: Rate</li> <li>■ Unit: Megabits per second (Mbps)</li> <li>■ Rollup Type: Average (Minimum/Maximum)</li> <li>■ Collection Level: 1 (4)</li> </ul>

## Chart Analysis

Network performance depends on the application workload and network configuration. Dropped network packets indicate a bottleneck in the network. To determine whether packets are being dropped, use `esxtop` or the advanced performance charts to examine the `droppedTx` and `droppedRx` network counter values.

If packets are being dropped, adjust the virtual machine shares. If packets are not being dropped, check the size of the network packets and the data receive and transfer rates. In general, the larger the network packets, the faster the network speed. When the packet size is large, fewer packets are transferred, which reduces the amount of CPU required to process the data. When network packets are small, more packets are transferred but the network speed is slower because more CPU is required to process the data.

---

**Note** In some instances, large packets might result in a high network latency. To check the network latency, use the VMware AppSpeed performance monitoring application or a third-party application.

---

If packets are not being dropped and the data receive rate is slow, the host is probably lacking the CPU resources required to handle the load. Check the number of virtual machines assigned to each physical NIC. If necessary, perform load balancing by moving virtual machines to different vSwitches or by adding more NICs to the host. You can also move virtual machines to another host or increase the host CPU or virtual machine CPU.

If you experience network-related performance problems, also consider taking the following actions.

**Table 1-118. Networking Performance Enhancement Advice**

#	Resolution
1	Verify that VMware Tools is installed on each virtual machine.
2	If possible, use vmxnet3 NIC drivers, which are available with VMware Tools. They are optimized for high performance.
3	If virtual machines running on the same host communicate with each other, connect them to the same vSwitch to avoid transferring packets over the physical network.
4	Assign each physical NIC to a port group and a vSwitch.
5	Use separate physical NICs to handle the different traffic streams, such as network packets generated by virtual machines, iSCSI protocols, vMotion tasks.
6	Ensure that the physical NIC capacity is large enough to handle the network traffic on that vSwitch. If the capacity is not enough, consider using a high-bandwidth physical NIC (10 Gbps). Alternatively, consider moving some virtual machines to a vSwitch with a lighter load or to a new vSwitch.
7	If packets are being dropped at the vSwitch port, increase the virtual network driver ring buffers where applicable.
8	Verify that the reported speed and duplex settings for the physical NIC match the hardware expectations and that the hardware is configured to run at its maximum capability. For example, verify that NICs with 1 Gbps are not reset to 100 Mbps because they are connected to an older switch.
9	Verify that all NICs are running in full duplex mode. Hardware connectivity problems might result in a NIC resetting itself to a lower speed or half duplex mode.
10	Use vNICs that are TCP Segmentation Offload (TSO)-capable, and verify that TSO-Jumbo Frames are enabled where possible.

## Network Rate (Mbps)

The Network Rate chart displays network usage for virtual machines.

This chart is located in the **Home** view of the Virtual Machine **Performance** tab. It appears only at collection levels 3 and 4.

**Table 1-119. Data Counters**

Chart Label	Description
Data Receive Rate	Rate at which data is received across each virtual NIC instance on the virtual machine. <ul style="list-style-type: none"> <li>■ Counter: received</li> <li>■ Stats Type: Rate</li> <li>■ Unit: Megabits per second (Mbps)</li> <li>■ Rollup Type: Average</li> <li>■ Collection Level: 2 (4)</li> </ul>
Data Transmit Rate	Rate at which data is transmitted across each virtual NIC instance on the virtual machine. <ul style="list-style-type: none"> <li>■ Counter: transmitted</li> <li>■ Stats Type: Rate</li> <li>■ Unit: Megabits per second (Mbps)</li> <li>■ Rollup Type: Average</li> <li>■ Collection Level: 2 (4)</li> </ul>

### Chart Analysis

Network performance depends on the application workload and network configuration. Dropped network packets indicate a bottleneck in the network. To determine whether packets are being dropped, use `esx top` or the advanced performance charts to examine the `droppedTx` and `droppedRx` network counter values.

If packets are being dropped, adjust the virtual machine shares. If packets are not being dropped, check the size of the network packets and the data receive and transfer rates. In general, the larger the network packets, the faster the network speed. When the packet size is large, fewer packets are transferred, which reduces the amount of CPU required to process the data. When network packets are small, more packets are transferred but the network speed is slower because more CPU is required to process the data.

---

**Note** In some instances, large packets might result in a high network latency. To check the network latency, use the VMware AppSpeed performance monitoring application or a third-party application.

---

If packets are not being dropped and the data receive rate is slow, the host is probably lacking the CPU resources required to handle the load. Check the number of virtual machines assigned to each physical NIC. If necessary, perform load balancing by moving virtual machines to different vSwitches or by adding more NICs to the host. You can also move virtual machines to another host or increase the host CPU or virtual machine CPU.

If you experience network-related performance problems, also consider taking the following actions.

**Table 1-120. Networking Performance Enhancement Advice**

#	Resolution
1	Verify that VMware Tools is installed on each virtual machine.
2	If possible, use <code>vmxnet3</code> NIC drivers, which are available with VMware Tools. They are optimized for high performance.
3	If virtual machines running on the same host communicate with each other, connect them to the same vSwitch to avoid transferring packets over the physical network.

**Table 1-120. Networking Performance Enhancement Advice (Continued)**

#	Resolution
4	Assign each physical NIC to a port group and a vSwitch.
5	Use separate physical NICs to handle the different traffic streams, such as network packets generated by virtual machines, iSCSI protocols, vMotion tasks.
6	Ensure that the physical NIC capacity is large enough to handle the network traffic on that vSwitch. If the capacity is not enough, consider using a high-bandwidth physical NIC (10 Gbps). Alternatively, consider moving some virtual machines to a vSwitch with a lighter load or to a new vSwitch.
7	If packets are being dropped at the vSwitch port, increase the virtual network driver ring buffers where applicable.
8	Verify that the reported speed and duplex settings for the physical NIC match the hardware expectations and that the hardware is configured to run at its maximum capability. For example, verify that NICs with 1 Gbps are not reset to 100 Mbps because they are connected to an older switch.
9	Verify that all NICs are running in full duplex mode. Hardware connectivity problems might result in a NIC resetting itself to a lower speed or half duplex mode.
10	Use vNICs that are TCP Segmentation Offload (TSO)-capable, and verify that TSO-Jumbo Frames are enabled where possible.

## Network Packets (Number)

The Network Packets monitors network bandwidth for virtual machines.

This chart is located in the **Home** view of the Virtual Machine **Performance** tab. It appears only at collection levels 3 and 4.

**Table 1-121. Data Counters**

Chart Label	Description
Packets Transmitted	<p>Number of network packets transmitted across the top ten virtual NIC instances on the virtual machine. The chart also displays the aggregated value for each NIC.</p> <ul style="list-style-type: none"> <li>■ Counter: packetTx</li> <li>■ Stats Type: Absolute</li> <li>■ Unit: Number</li> <li>■ Rollup Type: Summation</li> <li>■ Collection Level: 3</li> </ul>
Packets Received	<p>Number of network packets received across the top ten virtual NIC instances on the virtual machine. The chart also displays the aggregated value for each NIC.</p> <ul style="list-style-type: none"> <li>■ Counter: packetRx</li> <li>■ Stats Type: Absolute</li> <li>■ Unit: Number</li> <li>■ Rollup Type: Summation</li> <li>■ Collection Level: 3</li> </ul>

## Chart Analysis

Network performance depends on the application workload and network configuration. Dropped network packets indicate a bottleneck in the network. To determine whether packets are being dropped, use `esxtop` or the advanced performance charts to examine the `droppedTx` and `droppedRx` network counter values.

If packets are being dropped, adjust the virtual machine shares. If packets are not being dropped, check the size of the network packets and the data receive and transfer rates. In general, the larger the network packets, the faster the network speed. When the packet size is large, fewer packets are transferred, which reduces the amount of CPU required to process the data. When network packets are small, more packets are transferred but the network speed is slower because more CPU is required to process the data.

---

**Note** In some instances, large packets might result in a high network latency. To check the network latency, use the VMware AppSpeed performance monitoring application or a third-party application.

---

If packets are not being dropped and the data receive rate is slow, the host is probably lacking the CPU resources required to handle the load. Check the number of virtual machines assigned to each physical NIC. If necessary, perform load balancing by moving virtual machines to different vSwitches or by adding more NICs to the host. You can also move virtual machines to another host or increase the host CPU or virtual machine CPU.

If you experience network-related performance problems, also consider taking the following actions.

**Table 1-122. Networking Performance Enhancement Advice**

#	Resolution
1	Verify that VMware Tools is installed on each virtual machine.
2	If possible, use <code>vmxnet3</code> NIC drivers, which are available with VMware Tools. They are optimized for high performance.
3	If virtual machines running on the same host communicate with each other, connect them to the same vSwitch to avoid transferring packets over the physical network.
4	Assign each physical NIC to a port group and a vSwitch.
5	Use separate physical NICs to handle the different traffic streams, such as network packets generated by virtual machines, iSCSI protocols, vMotion tasks.
6	Ensure that the physical NIC capacity is large enough to handle the network traffic on that vSwitch. If the capacity is not enough, consider using a high-bandwidth physical NIC (10 Gbps). Alternatively, consider moving some virtual machines to a vSwitch with a lighter load or to a new vSwitch.
7	If packets are being dropped at the vSwitch port, increase the virtual network driver ring buffers where applicable.
8	Verify that the reported speed and duplex settings for the physical NIC match the hardware expectations and that the hardware is configured to run at its maximum capability. For example, verify that NICs with 1 Gbps are not reset to 100 Mbps because they are connected to an older switch.
9	Verify that all NICs are running in full duplex mode. Hardware connectivity problems might result in a NIC resetting itself to a lower speed or half duplex mode.
10	Use vNICs that are TCP Segmentation Offload (TSO)-capable, and verify that TSO-Jumbo Frames are enabled where possible.



## Space in GB

The Space in GB chart displays space utilization data counters for virtual machines.

This chart is located in the **Storage** view of the virtual machine **Performance** tab.

**Table 1-123. Data Counters**

Chart Label	Description
Allocated	<p>Total amount of logical datastore space provisioned by an administrator for the virtual machine. It is the storage size up to which the virtual machine files on datastores can grow. This includes log files, VMX files, and other miscellaneous files. Allocated space is not always in use.</p> <ul style="list-style-type: none"> <li>■ Counter: provisioned</li> <li>■ Stats Type: Absolute</li> <li>■ Unit: Gigabytes (GB)</li> <li>■ Rollup Type: Latest</li> <li>■ Collection Level: 1</li> </ul>
Used	<p>Amount of physical datastore space in use by the virtual machine files.</p> <ul style="list-style-type: none"> <li>■ Counter: used</li> <li>■ Stats Type: Absolute</li> <li>■ Unit: Gigabytes (GB)</li> <li>■ Rollup Type: Latest</li> <li>■ Collection Level: 1</li> </ul>
Not Shared	<p>Amount of datastore space that belongs only to this virtual machine and is not shared with other virtual machines. Only unshared space is guaranteed to be reclaimed for the virtual machine if, for example, it is moved to a different datastore and then back again. The value is an aggregate of all unshared space for the virtual machine, across all datastores.</p> <ul style="list-style-type: none"> <li>■ Counter: unshared</li> <li>■ Stats Type: Absolute</li> <li>■ Unit: Gigabytes (GB)</li> <li>■ Rollup Type: Latest</li> <li>■ Collection Level: 1</li> </ul>

### Chart Analysis

The datastore is at full capacity when the used space is equal to the capacity. Allocated space can be larger than datastore capacity, for example, when you have snapshots and thin-provisioned disks. If possible, you can provision more space to the datastore, or you can add disks to the datastore or use shared datastores.

If snapshot files are consuming high datastore space, consider consolidating them to the virtual disk when they are no longer needed. Consolidating the snapshots deletes the redo log files and removes the snapshots from the vSphere Web Client user interface. For information about consolidating the data center, see the vSphere documentation.

## Space Utilization by Datastores

The Space Utilization by Datastores chart displays the amount of space used by a virtual machine on different datastores in the data center.

**Note** This chart does not show historical statistics. It only shows the most recently available data, which may be up to 30 minutes late, depending on when the last statistics rollup occurred. In addition, statistics are not collected across all datastores at one time. They are collected asynchronously.

The Space Utilization by Datastores chart is located in the **Storage** view of the virtual machine **Performance** tab.

**Table 1-124. Data Counters**

Chart Label	Description
<i>datastore_name</i>	<p>Amount of disk space in the datastore currently in use by the virtual machine.</p> <ul style="list-style-type: none"> <li>■ Counter: used</li> <li>■ Stats Type: Absolute</li> <li>■ Unit: Gigabytes (GB)</li> <li>■ Rollup Type: Latest</li> <li>■ Collection Level: 1</li> </ul>

### Chart Analysis

The datastore is at full capacity when the used space is equal to the capacity. Allocated space can be larger than datastore capacity, for example, when you have snapshots and thin-provisioned disks. If possible, you can provision more space to the datastore, or you can add disks to the datastore or use shared datastores.

If snapshot files are consuming high datastore space, consider consolidating them to the virtual disk when they are no longer needed. Consolidating the snapshots deletes the redo log files and removes the snapshots from the vSphere Web Client user interface. For information about consolidating the data center, see the vSphere documentation.

## Space Utilization By File Type

The Space Utilization by File Type chart displays the datastore usage by virtual machine files.

**Note** This chart does not show historical statistics. It only shows the most recently available data, which may be up to 30 minutes late, depending on when the last statistics rollup occurred. In addition, statistics are not collected across all datastores at one time. They are collected asynchronously.

The Space Utilization by File Type chart is located in the **Storage** view of the virtual machine **Performance** tab.

## Datastore counters

**Table 1-125. Data Counters**

File Type	Description
Virtual Disks	<p>Amount of disk space used by virtual disk files.</p> <p>Virtual disk files store the contents of the virtual machine's hard disk drive, including information that you write to a virtual machine's hard disk - the operating system, program files, and data files. The files have the extension <code>.vmdk</code> and appear as a physical disk drive to a guest operating system.</p> <hr/> <p><b>Note</b> Delta disks, which also have an extension <code>.vmdk</code>, are not included in this file type.</p> <hr/> <ul style="list-style-type: none"> <li>■ Counter: used</li> <li>■ Stats Type: Absolute</li> <li>■ Unit: Gigabytes (GB)</li> <li>■ Rollup Type: Latest</li> <li>■ Collection Level: 1 (4)</li> </ul>
Swap Files	<p>Amount of disk space used by swap files.</p> <p>Swap files back up the virtual machine's physical memory.</p> <ul style="list-style-type: none"> <li>■ Counter: used</li> <li>■ Stats Type: Absolute</li> <li>■ Unit: Gigabytes (GB)</li> <li>■ Rollup Type: Latest</li> <li>■ Collection Level: 1 (4)</li> </ul>
Snapshots	<p>Amount of disk space used by virtual machine snapshot files.</p> <p>Snapshot files store information about virtual machine snapshots. They include snapshot state files and delta disk files. A snapshot state file stores the running state of the virtual machine at the time of the snapshot. It has the extension <code>.vmsn</code>. A delta disk file stores the updates made by the virtual machine to the virtual disks after a snapshot is taken.</p> <ul style="list-style-type: none"> <li>■ Counter: used</li> <li>■ Stats Type: Absolute</li> <li>■ Unit: GigaBytes (GB)</li> <li>■ Rollup Type: Latest</li> <li>■ Collection Level: 1 (4)</li> </ul>
Other VM Files	<p>Amount of disk space used by all other virtual machine files, such as configuration files and log files.</p> <ul style="list-style-type: none"> <li>■ Counter: used</li> <li>■ Stats Type: Absolute</li> <li>■ Unit: Gigabytes (GB)</li> <li>■ Rollup Type: Latest</li> <li>■ Collection Level: 1 (4)</li> </ul>
Total Space	<p>Amount of disk space used by the virtual machine.</p> <p>total space = virtual disk space + swap file space + snapshot space + other VM file space</p>

## Chart Analysis

The datastore is at full capacity when the used space is equal to the capacity. Allocated space can be larger than datastore capacity, for example, when you have snapshots and thin-provisioned disks. If possible, you can provision more space to the datastore, or you can add disks to the datastore or use shared datastores.

If snapshot files are consuming high datastore space, consider consolidating them to the virtual disk when they are no longer needed. Consolidating the snapshots deletes the redo log files and removes the snapshots from the vSphere Web Client user interface. For information about consolidating the data center, see the vSphere documentation.

## Fault Tolerance Performance Counters

The fault tolerance charts contain information about CPU and memory for fault tolerant virtual machines.

**Note** The performance charts and help topics for fault tolerance are available only when you have vSphere Fault Tolerance enabled. If you select a link for a secondary virtual machine in the thumbnail section of the Resource Pools and Virtual Machines view of the cluster Performance tab, the navigation in the inventory updates to the primary virtual machine. This occurs because secondary machines are not displayed in the inventory.

### CPU (MHz)

The CPU (MHz) chart displays virtual CPU usage for fault tolerant virtual machines.

This chart is located in the **Fault Tolerance** view of the virtual machine **Performance** tab. It is available only at collection levels 3 and 4.

**Table 1-126. Data Counters**

Name	Description
Usage	<p>The average amount of virtual CPU, per CPU instance, in use on the primary and secondary fault tolerant virtual machines.</p> <ul style="list-style-type: none"> <li>■ Counter: usagemhz</li> <li>■ Stats Type: Rate</li> <li>■ Unit: Megahertz (MHz)</li> <li>■ Rollup Type: Average (Minimum/Maximum)</li> <li>■ Collection Level: 3 (4)</li> </ul>

### Chart Analysis

A large discrepancy in CPU usage between the primary and secondary virtual machines might indicate performance problems. The CPU ready, system, and wait times of each virtual machine should be synchronized. A large discrepancy in these values might indicate performance problems. Consider taking the following actions.

**Table 1-127. CPU Performance Enhancement Advice**

#	Resolution
1	Verify that the primary and secondary hosts are in the same CPU model family and have similar CPU configurations. For best results, use CPUs with the same stepping level.
2	Verify that the CPU resource reservations set for both virtual machines are consistent within the cluster. VMware HA plans for a worst-case scenario by considering all powered on virtual machines in a cluster and finding the maximum memory and CPU reservations.
3	Verify that the network and datastore connections for both virtual machines are similar.

**Table 1-127. CPU Performance Enhancement Advice (Continued)**

#	Resolution
4	Turn off power management (also known as power-capping) in the BIOS. If power management is enabled, the secondary host might enter lower performance, power-saving modes. Such modes can leave the secondary virtual machine with insufficient CPU resources, potentially making it impossible for the secondary to complete all tasks completed on a primary in a timely fashion.
5	Turn off hyperthreading in the BIOS. If hyperthreading is enabled and the secondary virtual machine is sharing a CPU with another demanding virtual machine, the secondary virtual machine might run too slowly to complete all tasks completed on the primary in a timely fashion.

### CPU System Time for vCPU (%)

The CPU System Time chart displays virtual CPU usage for fault tolerant virtual machines.

This chart is located in the **Fault Tolerance** view of the Virtual Machine **Performance** tab. It is available only at collection levels 3 and 4.

**Table 1-128. Data Counters**

Chart Label	Description
System	Amount of time spent on system processes on each virtual CPU in the virtual machine.
	<b>Note</b> This is the host view of the CPU usage, not the guest operating system view.
	<ul style="list-style-type: none"> <li>■ Counter: system</li> <li>■ Stats Type: Delta</li> <li>■ Unit: Percentage (%)</li> <li>■ Rollup Type: Summation</li> <li>■ Collection Level: 3</li> </ul>

### Chart Analysis

A large discrepancy in CPU usage between the primary and secondary virtual machines might indicate performance problems. The CPU ready, system, and wait times of each virtual machine should be synchronized. A large discrepancy in these values might indicate performance problems. Consider taking the following actions.

**Table 1-129. CPU Performance Enhancement Advice**

#	Resolution
1	Verify that the primary and secondary hosts are in the same CPU model family and have similar CPU configurations. For best results, use CPUs with the same stepping level.
2	Verify that the CPU resource reservations set for both virtual machines are consistent within the cluster. VMware HA plans for a worst-case scenario by considering all powered on virtual machines in a cluster and finding the maximum memory and CPU reservations.
3	Verify that the network and datastore connections for both virtual machines are similar.

**Table 1-129. CPU Performance Enhancement Advice (Continued)**

#	Resolution
4	Turn off power management (also known as power-capping) in the BIOS. If power management is enabled, the secondary host might enter lower performance, power-saving modes. Such modes can leave the secondary virtual machine with insufficient CPU resources, potentially making it impossible for the secondary to complete all tasks completed on a primary in a timely fashion.
5	Turn off hyperthreading in the BIOS. If hyperthreading is enabled and the secondary virtual machine is sharing a CPU with another demanding virtual machine, the secondary virtual machine might run too slowly to complete all tasks completed on the primary in a timely fashion.

### CPU Used Time for vCPU (%)

The CPU Used Time chart displays virtual CPU usage for fault tolerant virtual machines.

This chart is located in the **Fault Tolerance** view of the Virtual Machine **Performance** tab. The chart is available only for collection levels 3 and 4.

**Table 1-130. Data Counters**

Chart Label	Description
used	<p>Amount of used virtual CPU as a percentage of total available CPU on the primary and secondary virtual machines.</p> <p>A high value indicates excessive use of CPU resources.</p> <ul style="list-style-type: none"> <li>■ Counter: used</li> <li>■ Stats Type: Delta</li> <li>■ Unit: Percentage (%)</li> <li>■ Rollup Type: Summation</li> <li>■ Collection Level: 1</li> </ul>

### Chart Analysis

A large discrepancy in CPU usage between the primary and secondary virtual machines might indicate performance problems. The CPU ready, system, and wait times of each virtual machine should be synchronized. A large discrepancy in these values might indicate performance problems. Consider taking the following actions.

**Table 1-131. CPU Performance Enhancement Advice**

#	Resolution
1	Verify that the primary and secondary hosts are in the same CPU model family and have similar CPU configurations. For best results, use CPUs with the same stepping level.
2	Verify that the CPU resource reservations set for both virtual machines are consistent within the cluster. VMware HA plans for a worst-case scenario by considering all powered on virtual machines in a cluster and finding the maximum memory and CPU reservations.
3	Verify that the network and datastore connections for both virtual machines are similar.

**Table 1-131. CPU Performance Enhancement Advice (Continued)**

#	Resolution
4	Turn off power management (also known as power-capping) in the BIOS. If power management is enabled, the secondary host might enter lower performance, power-saving modes. Such modes can leave the secondary virtual machine with insufficient CPU resources, potentially making it impossible for the secondary to complete all tasks completed on a primary in a timely fashion.
5	Turn off hyperthreading in the BIOS. If hyperthreading is enabled and the secondary virtual machine is sharing a CPU with another demanding virtual machine, the secondary virtual machine might run too slowly to complete all tasks completed on the primary in a timely fashion.

## Memory Active (MB)

The Memory Active chart displays active memory usage for fault tolerant virtual machines.

This chart is located in the **Fault Tolerance** view of the Virtual Machine **Performance** tab. It is not available at collection level 1.

**Table 1-132. Data Counters**

Chart Label	Description
Active	<p>Amount of guest physical memory in use by the fault tolerant virtual machine. Active memory is estimated by VMkernel statistical sampling and represents the actual amount of memory the virtual machine needs. Additional, unused memory may be swapped out or ballooned with no performance impact.</p> <ul style="list-style-type: none"> <li>■ Counter: active</li> <li>■ Stats Type: Absolute</li> <li>■ Unit: Megabytes (MB)</li> <li>■ Rollup Type: Average (Minimum/Maximum)</li> <li>■ Collection Level: 2 (4)</li> </ul> <p>Make sure that the primary and secondary virtual machines have enough memory. If the secondary system is not provisioned well, it might slow down performance of the primary virtual machine or fail.</p>

## Chart Analysis

A virtual machine's memory size must be slightly larger than the average guest memory usage. This enables the host to accommodate workload spikes without swapping memory among guests. Increasing the virtual machine memory size results in more overhead memory usage.

If sufficient swap space is available, a high balloon value does not cause performance problems. However, if the swapin and swapout values for the host are large, the host is probably lacking the amount of memory required to meet the demand.

If a virtual machine has high ballooning or swapping, check the amount of free physical memory on the host. A free memory value of 6% or less indicates that the host cannot meet the memory requirements. This leads to memory reclamation, which might degrade performance. If the active memory size is the same as the granted memory size, demand for memory is greater than the memory resources available. If the active memory is consistently low, the memory size might be too large.

If the host has enough free memory, check the resource shares, reservation, and limit of the virtual machines and resource pools on the host. Verify that the host settings are adequate and not lower than those set for the virtual machine.

If little free memory is available, or if you notice degradation in performance, consider taking the following actions.

**Table 1-133. Memory Performance Enhancement Advice**

#	Resolution
1	Verify that VMware Tools is installed on each virtual machine. The balloon driver is installed with VMware Tools and is critical to performance.
2	Verify that the balloon driver is enabled. The VMkernel regularly reclaims unused virtual machine memory by ballooning and swapping. Generally, this does not impact virtual machine performance.
3	Reduce the memory space on the virtual machine, and correct the cache size if it is too large. This frees up memory for other virtual machines.
4	If the memory reservation of the virtual machine is set to a value much higher than its active memory, decrease the reservation setting so that the VMkernel can reclaim the idle memory for other virtual machines on the host.
5	Migrate one or more virtual machines to a host in a DRS cluster.
6	Add physical memory to the host.

## Memory Swapout (MB)

The Memory Swapout chart displays the swapout memory usage for fault tolerant virtual machines.

This chart is located in the **Fault Tolerance** view of the Virtual Machine **Performance** tab. It is not available at collection level 1.

**Table 1-134. Data Counters**

Chart Label	Description
Swapout	<p>Amount of machine memory written to the VMkernel swap file.</p> <ul style="list-style-type: none"> <li>■ Counter: swapout</li> <li>■ Stats Type: Absolute</li> <li>■ Unit: MegaBytes</li> <li>■ Rollup Type: Average (Minimum/Maximum)</li> <li>■ Collection Level: 2 (4)</li> </ul> <p>Make sure that the primary and secondary virtual machines have enough memory and that the swapout value is not high. If the secondary system is not provisioned well, it might slow down performance of the primary virtual machine or fail.</p>

## Chart Analysis

A virtual machine's memory size must be slightly larger than the average guest memory usage. This enables the host to accommodate workload spikes without swapping memory among guests. Increasing the virtual machine memory size results in more overhead memory usage.

If there is sufficient swap space, a high balloon value is not a performance issue. However, if the swapin and swapout values for the host are large, the host is probably lacking the memory required to meet the demand.



If a virtual machine has high ballooning or swapping, check the amount of free physical memory on the host. The host might require more memory resources. If it does not, check the resource shares, reservation, and limit of the virtual machines and resource pools on the host. Verify that the host settings are adequate and not lower than those set for the virtual machine.

If memory usage is high or you notice degradation in performance, consider taking the following actions.

**Table 1-135. Memory Performance Enhancement Advice**

#	Resolution
1	Verify that VMware Tools is installed on each virtual machine. The balloon driver is installed with VMware Tools and is critical to performance.
2	Verify that the balloon driver is enabled. The VMkernel regularly reclaims unused virtual machine memory by ballooning and swapping. Generally, this does not impact virtual machine performance.
3	Reduce the memory space on the virtual machine, and correct the cache size if it is too large. This frees up memory for other virtual machines.
4	If the memory reservation of a virtual machine is set to a value much higher than its active memory, decrease the reservation setting so that the VMkernel can reclaim the idle memory for other virtual machines on the host.
5	Migrate one or more virtual machines to a host in a DRS cluster.
6	Add physical memory to the host.

## Working with Advanced and Custom Charts

Use advanced charts, or create your own custom charts, to see more performance data. Advanced charts can be useful when you are aware of a problem but need more statistical data to pinpoint the source of the trouble.

Advanced charts include the following features:

- More information. Hover over a data point in a chart and details about that specific data point are displayed.
- Customizable charts. Change chart settings. To create your own charts, save custom settings.
- Export to spreadsheet.
- Save to image file or spreadsheet.

## View Advanced Performance Charts

Advanced charts support data counters that are not supported in other performance charts.

### Procedure

- 1 Navigate to an inventory object in the vSphere Web Client.
- 2 Click the **Monitor** tab, and click **Performance**.
- 3 Click **Advanced**.

- 4 (Optional) To view a different chart, select an option from the **View** list.

The amount of historical data displayed in a chart depends on the collection interval and statistics level set for vCenter Server.

## Change Advanced Chart Settings

You can customize a performance chart by specifying the objects to monitor, the counters to include, the time range, and chart type. You can customize preconfigured chart views and create chart views.

### Procedure

- 1 Select an inventory object in the vSphere Web Client.
- 2 Click the **Monitor** tab, and click **Performance**.
- 3 Click **Advanced**.
- 4 Click **Chart Options**.
- 5 In Chart Metrics, select a metric group for the chart.
- 6 Select a time range for the metric group.

Time range options are not active unless you select **Custom interval** in the **Timespan** menu.

If you choose **Custom interval**, do one of the following.

- Select **Last** and set the number of hours, days, weeks, or months for the amount of time to monitor the object.
- Select **From** and select the beginning and end dates.

You can also customize the time range options by customizing the statistics collection interval setting.

- 7 In Target Objects, select the inventory objects to display in the chart.  
You can also specify the objects using the **All** or **None** buttons.
- 8 Select the chart type.

When selecting the stacked graph option, consider the following.

- You can select only one item from the list of measurements.
- Per-virtual-machine stacked graphs are available only for hosts.
- To display information about the counter's function and whether the selected metric can be stacked for per-virtual-machine graphs, click a counter description name.

- 9 In Counters, select the data counters to display in the chart.

You can also specify counters using the **All** or **None** buttons. The **All** button is inactive when there are more than two different counter units for the corresponding metric group.

- 10 Click **OK**.

## Create a Custom Advanced Chart

You can create your own charts by saving customized chart settings. New charts are added to the **View** menu and will appear there only when charts for the selected object are being displayed.

### Procedure

- 1 Select an inventory object in the vSphere Web Client.
- 2 Click the **Monitor** tab, click **Performance**, and navigate to the Chart Options dialog of a chart.
- 3 Customize chart settings.
- 4 Click **Save Options As...**
- 5 Enter a name for your settings.
- 6 Click **OK**.

The chart settings are saved and an entry for your chart is added to the **View** menu.

## Delete a Custom Advanced Chart View

You can delete custom chart views from the vSphere Web Client.

### Procedure

- 1 Select an inventory object in the vSphere Web Client.
- 2 Click the **Monitor** tab, and click **Performance**.
- 3 Click **Advanced**.
- 4 Click **Chart Options**.
- 5 Select a chart and click **Delete Options**.
- 6 Click **OK** to confirm deletion.

The chart is deleted, and it is removed from the **View** menu.

## Save Chart Data to a File

You can save data from the advanced performance charts to a file in various graphic formats or in comma-separated values (CSV) format.

### Procedure

- 1 In the vSphere Web Client, select an inventory object.
- 2 Click the **Monitor** tab, and click **Performance**.
- 3 Click **Advanced**.
- 4 Click the **Export** icon (↵).

## 5 Select a file type.

Option	Description
To PNG	Exports a bitmap image in the PNG format.
To JPEG	Exports a bitmap image in the JPEG format.
To CSV	Exports plain-text data in the CSV format.

## 6 Enter a name and location for the file.

## 7 Click **Save**.

The file is saved to the location and format you specified.

# Troubleshoot and Enhance Performance

This section presents tips for identifying and solving performance problems.

The suggestions in this section are not meant to be a comprehensive guide to diagnosing and troubleshooting problems in the virtual environment. It is meant to provide information about some common problems that can be solved without contacting VMware Technical Support.

## Solutions for Consistently High CPU Usage

Temporary spikes in CPU usage indicate that you are making the best use of CPU resources. Consistently high CPU usage might indicate a problem. You can use the CPU performance charts to monitor CPU usage for hosts, clusters, resource pools, virtual machines, and vApps.

### Problem

- Host CPU usage constantly is high. A high CPU usage value can lead to increased ready time and processor queuing of the virtual machines on the host.
- Virtual machine CPU usage is above 90% and the CPU ready value is above 20%. Application performance is impacted.

### Cause

- The host probably is lacking the CPU resources required to meet the demand.
- There might be too many virtual CPUs relative to the number of regular CPUs.
- There might be an IO storage or networking operation that places the CPU in a wait state.
- The Guest OS generates too much load for the CPU.

### Solution

- Verify that VMware Tools is installed on every virtual machine on the host.
- Compare the CPU usage value of a virtual machine with the CPU usage of other virtual machines on the host or in the resource pool. The stacked bar chart on the host's **Virtual Machine** view shows the CPU usage for all virtual machines on the host.

- Determine whether the high ready time for the virtual machine resulted from its CPU usage time reaching the CPU limit setting. If so, increase the CPU limit on the virtual machine.
- Increase the CPU shares to give the virtual machine more opportunities to run. The total ready time on the host might remain at the same level if the host system is constrained by CPU. If the host ready time does not decrease, set the CPU reservations for high-priority virtual machines to guarantee that they receive the required CPU cycles.
- Increase the amount of memory allocated to the virtual machine. This action decreases disk and or network activity for applications that cache. This might lower disk I/O and reduce the need for the host to virtualize the hardware. Virtual machines with smaller resource allocations generally accumulate more CPU ready time.
- Reduce the number of virtual CPUs on a virtual machine to only the number required to execute the workload. For example, a single-threaded application on a four-way virtual machine only benefits from a single vCPU. But the hypervisor's maintenance of the three idle vCPUs takes CPU cycles that could be used for other work.
- If the host is not already in a DRS cluster, add it to one. If the host is in a DRS cluster, increase the number of hosts and migrate one or more virtual machines onto the new host.
- Upgrade the physical CPUs or cores on the host if necessary.
- Use the newest version of hypervisor software, and enable CPU-saving features such as TCP Segmentation Offload, large memory pages, and jumbo frames.

## Solutions for Memory Performance Problems

Host machine memory is the hardware backing for guest virtual memory and guest physical memory. Host machine memory must be at least slightly larger than the combined active memory of the virtual machines on the host. A virtual machine's memory size must be slightly larger than the average guest memory usage. Increasing the virtual machine memory size results in more overhead memory usage.

### Problem

- Memory usage is constantly high (94% or greater) or constantly low (24% or less).
- Free memory consistently is 6% or less and swapping frequently occurs.

### Cause

- The host probably is lacking the memory required to meet the demand. The active memory size is the same as the granted memory size, which results in memory resources that are not sufficient for the workload. Granted memory is too much if the active memory is constantly low.
- Host machine memory resources are not enough to meet the demand, which leads to memory reclamation and degraded performance.
- The active memory size is the same as the granted memory size, which results in memory resources that are not sufficient for the workload.

**Solution**

- Verify that VMware Tools is installed on each virtual machine. The balloon driver is installed with VMware Tools and is critical to performance.
- Verify that the balloon driver is enabled. The VMkernel regularly reclaims unused virtual machine memory by ballooning and swapping. Generally, this does not impact virtual machine performance.
- Reduce the memory space on the virtual machine, and correct the cache size if it is too large. This frees up memory for other virtual machines.
- If the memory reservation of the virtual machine is set to a value much higher than its active memory, decrease the reservation setting so that the VMkernel can reclaim the idle memory for other virtual machines on the host.
- Migrate one or more virtual machines to a host in a DRS cluster.
- Add physical memory to the host.

**Solutions for Storage Performance Problems**

Datstores represent storage locations for virtual machine files. A storage location can be a VMFS volume, a directory on Network Attached Storage, or a local file system path. Datstores are platform-independent and host-independent.

**Problem**

- Snapshot files are consuming a lot of datastore space.
- The datastore is at full capacity when the used space is equal to the capacity. Allocated space can be larger than datastore capacity, for example, when you have snapshots and thin-provisioned disks.

**Solution**

- Consider consolidating snapshots to the virtual disk when they are no longer needed. Consolidating the snapshots deletes the redo log files and removes the snapshots from the vSphere Web Client user interface.
- You can provision more space to the datastore if possible, or you can add disks to the datastore or use shared datastores.

**Solutions for Disk Performance Problems**

Use the disk charts to monitor average disk loads and to determine trends in disk usage. For example, you might notice a performance degradation with applications that frequently read from and write to the hard disk. If you see a spike in the number of disk read/write requests, check if any such applications were running then.

**Problem**

- The value for the kernelLatency data counter is greater than 4 ms.

- The value for the deviceLatency data counter is greater than 15 ms indicates that there are probably problems with the storage array.
- The queueLatency data counter measures above zero.
- Spikes in latency.
- Unusual increases in read/write requests.

#### Cause

- The virtual machines on the host are trying to send more throughput to the storage system than the configuration supports.
- The storage array probably is experiencing internal problems.
- The workload is too high and the array cannot process the data fast enough.

#### Solution

- The virtual machines on the host are trying to send more throughput to the storage system than the configuration supports. Check the CPU usage, and increase the queue depth.
- Move the active VMDK to a volume with more spindles or add disks to the LUN.
- Increase the virtual machine memory. It should allow for more operating system caching, which can reduce I/O activity. Note: It may require you to increase the host memory. Increasing memory might reduce the need to store data because databases can utilize system memory to cache data and avoid disk access.
- Check swap statistics in the guest operating system to verify that virtual machines have adequate memory. Increase the guest memory, but not to an extent that leads to excessive host memory swapping. Install VMware Tools so that memory ballooning can occur.
- Defragment the file systems on all guests.
- Disable antivirus on-demand scans on the VMDK and VMEM files.
- Use the vendor's array tools to determine the array performance statistics. When too many servers simultaneously access common elements on an array, the disks might have trouble keeping up. Consider array-side improvements to increase throughput.
- Use Storage vMotion to migrate I/O-intensive virtual machines across multiple hosts.
- Balance the disk load across all physical resources available. Spread heavily used storage across LUNs that are accessed by different adapters. Use separate queues for each adapter to improve disk efficiency.
- Configure the HBAs and RAID controllers for optimal use. Verify that the queue depths and cache settings on the RAID controllers are adequate. If not, increase the number of outstanding disk requests for the virtual machine by adjusting the `Disk.SchedNumReqOutstanding` parameter.
- For resource-intensive virtual machines, separate the virtual machine's physical disk drive from the drive with the system page file. This alleviates disk spindle contention during periods of high use.

- On systems with sizable RAM, disable memory trimming by adding the line `MemTrimRate=0` to the virtual machine's VMX file.
- If the combined disk I/O is higher than a single HBA capacity, use multipathing or multiple links.
- For ESXi hosts, create virtual disks as preallocated. When you create a virtual disk for a guest operating system, select **Allocate all disk space now**. The performance degradation associated with reassigning additional disk space does not occur, and the disk is less likely to become fragmented.
- Use the most current hypervisor software.

## Solutions for Poor Network Performance

Network performance is dependent on application workload and network configuration. Dropped network packets indicate a bottleneck in the network. Slow network performance can be a sign of load-balancing problems.

### Problem

Network problems can manifest in many ways:

- Packets are being dropped.
- Network latency is high.
- Data receive rate is low.

### Cause

Network problems can have several causes:

- Virtual machine network resource shares are too few.
- Network packet size is too large, which results in high network latency. Use the VMware AppSpeed performance monitoring application or a third-party application to check network latency.
- Network packet size is too small, which increases the demand for the CPU resources needed for processing each packet. Host CPU, or possibly virtual machine CPU, resources are not enough to handle the load.

### Solution

- Determine whether packets are being dropped by using `esxtop` or the advanced performance charts to examine the `droppedTx` and `droppedRx` network counter values. Verify that VMware Tools is installed on each virtual machine.
- Check the number of virtual machines assigned to each physical NIC. If necessary, perform load balancing by moving virtual machines to different virtual switches or by adding more NICs to the host. You can also move virtual machines to another host or increase the host CPU or virtual machine CPU.
- If possible, use `vmxnet3` NIC drivers, which are available with VMware Tools. They are optimized for high performance.



- If virtual machines running on the same host communicate with each other, connect them to the same virtual switch to avoid the cost of transferring packets over the physical network.
- Assign each physical NIC to a port group and a virtual switch.
- Use separate physical NICs to handle the different traffic streams, such as network packets generated by virtual machines, iSCSI protocols, vMotion tasks.
- Ensure that the physical NIC capacity is large enough to handle the network traffic on that virtual switch. If the capacity is not enough, consider using a high-bandwidth physical NIC (10Gbps) or moving some virtual machines to a virtual switch with a lighter load or to a new virtual switch.
- If packets are being dropped at the virtual switch port, increase the virtual network driver ring buffers where applicable.
- Verify that the reported speed and duplex settings for the physical NIC match the hardware expectations and that the hardware is configured to run at its maximum capability. For example, verify that NICs with 1Gbps are not reset to 100Mbps because they are connected to an older switch.
- Verify that all NICs are running in full duplex mode. Hardware connectivity issues might result in a NIC resetting itself to a lower speed or half duplex mode.
- Use vNICs that are TCP Segmentation Offload (TSO)-capable, and verify that TCP Segmentation Offload-Jumbo Frames are enabled where possible.

## Empty Performance Charts

No graphics or data are displayed in performance charts.

### Problem

When data is missing for a performance chart, the chart appears empty and you see the message `No data available`.

### Cause

The causes for missing data in performance charts that are described herein are based on the assumption that the default roll-up configuration for the vCenter Server system has not changed. The causes include but are not limited to the following scenarios:

- Metrics introduced in ESXi 5.0 are not available for hosts running earlier versions.
- Data is deleted when you remove or add objects to vCenter Server.
- Performance charts data for inventory objects that were moved to a new site by VMware vCenter Site Recovery Manager is deleted from the old site and not copied to the new site.
- Performance charts data is deleted when you use VMware vMotion across vCenter Server instances.
- Real-time statistics are not available for disconnected hosts or powered off virtual machines.

- Real-time statistics are collected on hosts and aggregated every 5 minutes. After six data points are collected for approximately 30 minutes, they are rolled up to the vCenter Server database to create the 1-Day statistic. 1-Day statistics might not be available for 30 minutes after the current time, depending on when the sample period began.
- The 1-Day statistics are rolled up to create one data point every 30 minutes. If a delay occurs in the roll-up operation, the 1-Week statistics might not be available for 1 hour after the current time. It takes 30 minutes for the 1-Week collection interval, plus 30 minutes for the 1-Day collection interval.
- The 1-Week statistics are rolled up to create one data point every two hours. If a delay occurs in the roll-up operations, the 1-Month statistics might not be available for 3 hours. It takes 2 hours for the 1-Month collection interval, plus 1 hour for the 1-Week collection interval.
- The 1-Month statistics are rolled up to create one data point every day. If a delay occurs in the roll-up operations, the statistics might not be available for 1-day and 3 hours. It takes one day for the past year collection interval, plus 3 hours for the past month collection interval. During this time, the charts are empty.

#### **Solution**

- ◆ There is no solution available.

# Monitoring Guest Operating System Performance

# 2

This section describes how to install and view VMware-specific performance data for virtual machines that run Microsoft Windows operating systems. VMware provides performance counters that enable you to view data on many aspects of guest operating system performance for the Microsoft Windows Perfmon utility.

Some virtualization processes dynamically allocate available resources depending on the status, or utilization rates, of virtual machines in the environment. It can make obtaining accurate information about the resource utilization (CPU utilization, in particular) of individual virtual machines, or applications running within virtual machines, difficult. VMware now provides virtual machine-specific performance counter libraries for the Windows Perfmon utility. It enables application administrators to access accurate virtual machine resource utilization statistics from within the Windows Perfmon utility.

You can take advantage of virtualized CPU performance counters to use performance tuning tools inside the guest operating system. See the *vSphere Virtual Machine Administration* documentation.

This chapter includes the following topics:

- [Enable Statistics Collection for Guest Operating System Performance Analysis](#)
- [View Performance Statistics for Windows Guest Operating Systems](#)

## Enable Statistics Collection for Guest Operating System Performance Analysis

VMware-specific performance objects are loaded into Microsoft Windows Perfmon and enabled when VMware Tools is installed.

To display a performance chart for any performance object, you must add counters. See [View Performance Statistics for Windows Guest Operating Systems](#)

## View Performance Statistics for Windows Guest Operating Systems

You can display VMware specific statistics in the Microsoft Windows Perfmon utility.

### Prerequisites

Verify that a virtual machine with a Microsoft Windows operating system and VMware Tools is installed.

## Procedure

- 1 Open a console to the virtual machine and log in.
- 2 Select **Start > Run**.
- 3 Enter **Perfmon** and press **Enter**.
- 4 In the Performance dialog box, click **Add** .
- 5 In the Add Counters dialog box, select **Use local computer counters**.
- 6 Select a virtual machine performance object.  
Virtual machine performance object names begin with **VM**.
- 7 Select the counters that you want to display for that object.
- 8 If the performance object has multiple instances, select the instances you want to display.
- 9 Click **Add**.  
The Performance dialog box displays data for the selected performance object.
- 10 Click **Close** to close the Add Counter dialog box and return to the Performance dialog box.

## Monitoring Host Health Status

You can use the vSphere Web Client or the vSphere Client to monitor the state of host hardware components, such as CPU processors, memory, fans, and other components.

The host health monitoring tool allows you to monitor the health of a variety of host hardware components including:

- CPU processors
- Memory
- Fans
- Temperature
- Voltage
- Power
- Network
- Battery
- Storage
- Cable/Interconnect
- Software components
- Watchdog
- PCI devices
- Other

The host health monitoring tool presents data gathered using Systems Management Architecture for Server Hardware (SMASH) profiles. The information displayed depends on the sensors available on your server hardware. SMASH is an industry standard specification providing protocols for managing a variety of systems in the data center. For more information, see <http://www.dmtf.org/standards/smash>.

You can monitor host health status either by connecting the vSphere Web Client or the vSphere Client to a vCenter Server system. You can also set alarms to trigger when the host health status changes.

---

**Note** The interpretation of hardware monitoring information is specific for each hardware vendor. Your hardware vendor can help you understand the results of the host hardware components monitoring.

---

This chapter includes the following topics:

- [Monitor Health Status in the vSphere Client](#)
- [Monitor Hardware Health Status](#)
- [Reset Health Status Sensors in the vSphere Client](#)
- [Reset Health Status Sensors](#)

## Monitor Health Status in the vSphere Client

You can monitor the health status of host hardware in the vSphere Client

### Procedure

- 1 Select a host in the vSphere Client
- 2 Click the **Monitor** tab, and click **Hardware Health**
- 3 Select the type of information to view.

Option	Description
<b>Sensors</b>	Displays all sensors arranged in a tree view. If the status is blank, the health monitoring service cannot determine the status of the component.
<b>Alerts and warnings</b>	Displays alerts and warnings.
<b>System event log</b>	Displays the system event log.

## Monitor Hardware Health Status

You can monitor the health status of host hardware in the vSphere Web Client.

### Procedure

- 1 Select a host in the vSphere Web Client navigator.
- 2 Click the **Monitor** tab, and click **Hardware Status**.
- 3 Select the type of information to view.

Option	Description
<b>Sensors</b>	<p>Displays all sensors arranged in a tree view. If the status is blank, the health monitoring service cannot determine the status of the component.</p> <ul style="list-style-type: none"> <li>■ Click the <b>Expand All</b> icon to expand the tree view to show all sensors under each group.</li> <li>■ Click <b>Collapse All</b> icon to expand the tree view to show descriptive details for every sensor.</li> </ul>
<b>Alerts and warnings</b>	Displays alerts and warnings.
<b>System event log</b>	Displays the system event log.

## Reset Health Status Sensors in the vSphere Client

Some host hardware sensors display data that is cumulative over time. You can reset these sensors to clear the data in them and begin collecting new data.

If you need to preserve sensor data for troubleshooting or other purposes, take a screenshot, export the data, or download a support bundle before resetting sensors.

### Prerequisites

Verify that the vCenter Hardware Status plug-in is enabled.

### Procedure

- 1 Select a host in the vSphere Client
- 2 Click the **Monitor** tab, and click **Hardware Health**
- 3 Click **Reset sensors**.

## Reset Health Status Sensors

Some host hardware sensors display data that is cumulative over time. You can reset these sensors to clear the data in them and begin collecting new data.

If you need to preserve sensor data for troubleshooting or other purposes, take a screenshot, export the data, or download a support bundle before resetting sensors.

### Prerequisites

Verify that the vCenter Hardware Status plug-in is enabled.

### Procedure

- 1 Select a host in the vSphere Web Client navigator.
- 2 Click the **Monitor** tab, and click **Hardware Status**.
- 3 Click **Reset sensors**.

# Monitoring vSphere Health

You can check the health of vSphere Host and vCenter Server.

## Check vSphere Health in vSphere Client

You can use the vSphere Online health checks to monitor the health of the system. You can run health checks and send the data to VMware for advanced analysis.

### Prerequisites

- You must participate in the Customer Experience Improvement Program to use online health checks.
- To perform the online health checks, vCenter Server must be able to communicate over the Internet.

---

**Note** If Customer Improvement Experience Program (CEIP) is not enabled, the Internet connectivity check is unavailable.

---

### Procedure

- 1 Navigate to vCenter Server or select a host in the vSphere Client navigator
- 2 Click the **Monitor** tab, and click **Health**
- 3 Expand **Online health** to check the categories
- 4 Select the type of information to view

Option	Description
<b>Customer experience improvement program (CEIP)</b>	The CEIP check verifies whether the program is enabled for your vCenter Server. If it is not, click the button next to the health check, navigate to the CEIP page, and enroll in the program. To enable CEIP, click <b>Configure CEIP</b>
<b>Online health connectivity</b>	The Internet connectivity check verifies that vCenter Server can communicate with <i>vmware.com</i> over the HTTPS/443 interface. If communication is successful, this check passes. If communication fails, the check indicates that the Internet connectivity is not available.

vSphere periodically retests the health check and updates the results.



- 5 To run the health checks and update the results immediately, click the **Retest** button.

You can click the **Ask VMware** button to open a knowledge base article that describes the health check and provides information about how to resolve the issue.

# Monitoring Events, Alarms, and Automated Actions

# 5

vSphere includes a user-configurable events and alarms subsystem. This subsystem tracks events happening throughout vSphere and stores the data in log files and the vCenter Server database. This subsystem also enables you to specify the conditions under which alarms are triggered. Alarms can change state from mild warnings to more serious alerts as system conditions change, and can trigger automated alarm actions. This functionality is useful when you want to be informed, or take immediate action, when certain events or conditions occur for a specific inventory object, or group of objects.

## Events

Events are records of user actions or system actions that occur on objects in vCenter Server or on a host. Actions that might be recorded as events include, but are not limited to, the following examples:

- A license key expires
- A virtual machine is powered on
- A user logs in to a virtual machine
- A host connection is lost

Event data includes details about the event such as who generated it, when it occurred, and what type of event it is. There are three types of events:

- Information
- Warning
- Error

In the vSphere Web Client, event data is displayed in the **Monitor** tab. See [View Events](#).

## Alarms

Alarms are notifications that are activated in response to an event, a set of conditions, or the state of an inventory object. An alarm definition consists of the following elements in the vSphere Client:

- Name and description - Provides an identifying label and description.
- Targets - Defines the type of object that is monitored.

- Alarm Rules - Defines the event, condition, or state that triggers the alarm and defines the notification severity. It also defines operations that occur in response to triggered alarms.
- Last modified - The last modified date and time of the defined alarm.

An alarm definition consists of the following elements in the vSphere Web Client:

- Name and description - Provides an identifying label and description.
- Alarm type - Defines the type of object that is monitored.
- Triggers - Defines the event, condition, or state that triggers the alarm and defines the notification severity.
- Tolerance thresholds (Reporting) - Provides additional restrictions on condition and state triggers thresholds that must be exceeded before the alarm is triggered. Thresholds are not available in the vSphere Web Client.
- Actions - Defines operations that occur in response to triggered alarms. VMware provides sets of predefined actions that are specific to inventory object types.

Alarms have the following severity levels:

- Normal – green
- Warning – yellow
- Alert – red

Alarm definitions are associated with the object selected in the inventory. An alarm monitors the type of inventory objects specified in its definition.

For example, you might want to monitor the CPU usage of all virtual machines in a specific host cluster. You can select the cluster in the inventory, and add a virtual machine alarm to it. When enabled, that alarm monitors all virtual machines running in the cluster and triggers when any one of them meets the criteria defined in the alarm. To monitor a specific virtual machine in the cluster, but not others, select that virtual machine in the inventory and add an alarm to it. To apply the same alarms to a group of objects, place those objects in a folder and define the alarm on the folder.

---

**Note** You can enable, disable, and modify alarms only from the object in which the alarm is defined. For example, if you defined an alarm in a cluster to monitor virtual machines, you can only enable, disable, or modify that alarm through the cluster. You cannot change the alarm at the individual virtual machine level.

---

## Alarm Actions

Alarm actions are operations that occur in response to the trigger. For example, you can have an email notification sent to one or more administrators when an alarm is triggered.

---

**Note** Default alarms are not preconfigured with actions. You must manually set what action occurs when the triggering event, condition, or state occurs.

---

This chapter includes the following topics:

- [View Events](#)
- [View System Logs](#)
- [Export Events Data](#)
- [Streaming Events to a Remote Syslog Server](#)
- [Retention of Events in the vCenter Server Database](#)
- [View Triggered Alarms and Alarm Definitions](#)
- [Live Refresh of Recent Tasks and Alarms](#)
- [Set an Alarm in the vSphere Web Client](#)
- [Set an Alarm in the vSphere Client](#)
- [Acknowledge Triggered Alarms](#)
- [Reset Triggered Event Alarms](#)
- [Preconfigured vSphere Alarms](#)

## View Events

You can view events associated with a single object or view all vSphere events. The events list for a selected inventory object includes events associated with child objects. vSphere keeps information about tasks and events. It is set to a default period of 30 days and it is configurable.

### Procedure

- 1 Select an inventory object in the vSphere Web Client.
- 2 Click the **Monitor** tab, and click **Events**.
- 3 To see details, select an event.
- 4 (Optional) To filter the list, use the filter controls above the list.
- 5 (Optional) Click a column heading to sort the list.

## View System Logs

vSphere records events in the vCenter Server database. System log entries include such information as who generated the event, when the event was created, and the type of event.

### Prerequisites

- Required privilege: **Global.Diagnostics**

### Procedure

- 1 Select a vCenter Server instance in the vSphere Web Client navigator.

- 2 Click **Monitor**, and click **System Logs**.
- 3 From the drop-down menu, select the log.
- 4 (Optional) Click **Show All Lines** or **Show Next 2000 Lines** to see additional log entries.

## Export Events Data

You can export all or part of the events data stored in the vCenter Server database.

### Prerequisites

Required Role: **Read-only**

### Procedure

- 1 Select an inventory object in the vSphere Web Client.
- 2 Click the **Monitor** tab, and click **Events**.
- 3 Click the **Export** icon (📄).
- 4 In the **Export Events** window, specify what types of event information you want to export.
- 5 Click **Generate CSV Report**, and click **Save**.
- 6 Specify a file name and location and save the file.

## Streaming Events to a Remote Syslog Server

After you enable remote streaming, vCenter Server Appliance starts streaming and only the newly generated events are streamed to the remote syslog server.

All syslog messages begin with a specific prefix. You can distinguish the vCenter Server Appliance events from other syslog messages by their Event prefix.

The syslog protocol limits the length of syslog messages to 1024 characters. Messages that are longer than 1024 characters split into multiple syslog messages.

In the syslog server, events have the following format:

```
<syslog-prefix> : Event [eventId] [partInfo] [createdTime] [eventType] [severity] [user] [target]
[chainId] [desc]
```

Item	Description
syslog-prefix	Displays the syslog prefix. The <syslog-prefix> is determined by the remote syslog server configuration.
eventId	Displays the unique ID of the event message. The default value is Event.
partInfo	Displays whether the message is split into parts.
createdTime	Displays the time when the event was generated.
eventType	Displays the event type.
severity	Displays whether the event is a piece information, a warning, or an error.

Item	Description
user	Displays the name of the user who generated the event.
target	Displays the object the event refers to.
chainId	Displays information about the parent or the group ID.
desc	Displays the description of the event.

## Example: Split of Long Event Message into Multiple Syslog Messages

Events that are longer than 1024 characters split into multiple syslog messages in the following manner:

```
<syslog-prefix> : Event [eventId] [1-X] [payload-part-1]
<syslog-prefix> : Event [eventId] [2-X] [payload-part-2]
...
<syslog-prefix> : Event [eventId] [X-X] [payload-part-X]
```

The X stands for the number of the event message parts.

## Redirect vCenter Server Appliance Log Files to Another Machine

You can redirect the vCenter Server Appliance log files to another machine, for example, if you want to preserve storage space on the vCenter Server Appliance.

### Prerequisites

Log in to the vCenter Server Appliance Management Interface as root.

### Procedure

- 1 In the vCenter Server Appliance Management Interface, select **Syslog Configuration**.
- 2 Click **Edit**.
- 3 From the **Common Log Level** drop-down menu, select the log files to redirect.

Option	Description
*	All log files are redirected to the remote machine.
info	Only informational log files are redirected to the remote machine.
notice	Only notices are redirected to the remote machine. Notice indicates normal but significant condition.
warn	Only warnings are redirected to the remote machine.
error	Only error messages are redirected to the remote machine.
crit	Only critical log files are redirected to the remote machine.
alert	Only alerts are redirected to the remote machine. Alert indicates that action must be taken immediately.
emerg	Only emergency log files are redirected to the remote machine. Emergency indicates that the system stopped responding and cannot be used.

- 4 In the **Remote Syslog Host** text box, enter the FQDN or IP address of the machine on which you want to export the log files.
- 5 In the **Remote Syslog Port** text box, enter the port number to use for communication with the machine on which you want to export the log files.
- 6 From the **Remote Syslog Protocol** drop-down menu, select the protocol to use.

Option	Description
TCP	Transmission Control Protocol
UDP	User Datagram Protocol
TLS	Transport Layer Security
RELP	Reliable Event Logging Protocol

- 7 Click **OK**.

The new configuration settings are shown in the Remote Syslog Configuration pane.

- 8 (Optional) To stop redirecting log files to another machine, click **Reset**.

## Configure Streaming of Events to a Remote Syslog Server

You can also configure writing of events to the vCenter Server Appliance streaming facility. Streaming events is supported only for the vCenter Server Appliance. The streaming of events to a remote syslog server is disabled by default. You can enable and configure the streaming of vCenter Server events to a remote syslog server from the vCenter Server Appliance Management Interface.

### Procedure

- 1 In the vSphere Client, navigate to the vCenter Server instance.
- 2 Select the **Configure** tab.
- 3 Under **Settings**, Select **Advanced Settings**.
- 4 Click **Edit**.
- 5 In the **Filter** text box, type **vpxd.event**, and press Enter.
- 6 Enable or disable the `vpxd.event.syslog` option.

The default value for the setting is enabled.

## Retention of Events in the vCenter Server Database

You can configure vCenter Server to retain events in the database for a limited period. Discarding events periodically ensures optimal performance of the database.

In new installations of vCenter Server 6.5, the event clean up option is enabled by default and the default number of days to retain event messages in the database is 30. You can change this value to the number of days that you want to retain the event messages in the database.

If you are upgrading or migrating from vCenter Server 6.5 or earlier, and you had the event cleanup option enabled, your setting to retain events is preserved after the upgrade or migration to vCenter Server Appliance 6.5.

After the retention period ends, the events are deleted from the database. However, there might be latency in the deletion of the events that are older than the configured retention setting.

## Configure Database Settings

You can configure the maximum number of database connections that can occur simultaneously. To limit the growth of the vCenter Server database and save storage space, you can configure the database to discard information about tasks or events periodically.

---

**Note** Do not use the database retention options if you want to keep a complete history of tasks and events for your vCenter Server.

---

### Procedure

- 1 In the vSphere Client, navigate to the vCenter Server instance.
- 2 Select the **Configure** tab.
- 3 Under **Settings**, select **General**.
- 4 Click **Edit**.
- 5 Select **Database**.
- 6 In **Maximum connections**, type a number.  
  
Increase this number if your vCenter Server system performs many operations frequently and performance is critical. Decrease this number if the database is shared and connections to the database are costly. Do not change this value unless one of these issues pertains to your system.
- 7 Select the **Enabled** check box next to Task cleanup to have vCenter Server periodically delete the retained tasks.
- 8 (Optional) In **Tasks retained for**, type a value in days.  
  
Information about tasks that are performed on this vCenter Server system is discarded after the specified number of days.
- 9 Select the **Enabled** check box next to Event cleanup to have vCenter Server periodically clean up the retained events.
- 10 (Optional) In **Events retention**, type a value in days.  
  
Information about events for this vCenter Server system is discarded after the specified number of days.
- 11 Click **OK**.



## View Triggered Alarms and Alarm Definitions

Triggered alarms are visible in several locations throughout the vSphere Web Client.

### Procedure

- To view all triggered alarms, click **All** in the Alarms sidebar panel.

---

**Note** The list of alarms in the sidebar refreshes every 120 seconds. For information about changing the default refresh period, see the VMware knowledge base article at <http://kb.vmware.com/kb/2020290>.

---

- To view only newly triggered alarms, click **New** in the Alarms sidebar panel.  
The sidebar panel displays the latest 30 most critical alarms.
- To view acknowledged alarms, click **Acknowledged** in the Alarms sidebar panel.
- To view alarms triggered on a selected inventory object, click the **Monitor** tab, click **Issues**, and click **Triggered Alarms**.
- To view a list of available alarm definitions for a selected inventory object, click the **Monitor** tab, click **Issues**, and click **Alarm Definitions**.

## Live Refresh of Recent Tasks and Alarms

You can configure the vSphere Web Client to live refresh the recent tasks and the alarms that result from operations that other users perform in your environment.

By design the vSphere Web Client displays tasks initiated by other users and the resulting alarms from these tasks only when you manually refresh the vSphere Web Client. If you want to see the tasks from other users, or monitor alarms resulting from other users actions, perform the following procedure.

### Procedure

- 1 On the computer where the vSphere Web Client is installed, locate the `webclient.properties` file.

The location of this file depends on the operating system on which the vSphere Web Client is installed.

Operating System	File path
Windows	C:\ProgramData\VMware\vCenterServer\cfg\vsphere-client\webclient.properties
vCenter Server Appliance	/etc/vmware/vsphere-client/webclient.properties

- 2 Open the `webclient.properties` file, add the following configuration line, and save it.

```
live.updates.enabled=true
```

Live refresh of recent tasks and alarms is enabled for the vSphere Web Client.

- 3 Log out from the vSphere Web Client.
- 4 Use `https://hostname:9443/vsphere-client/` to log in to the vSphere Web Client.

*hostname* stands for the name or the IP address of the host where vCenter Server system runs.

If you log in to the vSphere Web Client by using the `https://hostname/vsphere-client/`, you will see no recent tasks or alarms under the respective Recent Tasks or Alarms portlets in the vSphere Web Client.

In an environment with multiple vCenter Server systems that are connected to the same vCenter Server Single-Sign On domain, the vSphere Web Client that you configured for live refresh displays recent tasks and alarms for all the vCenter Server instances in the domain. However, if you log in to a different vSphere Web Client, you will not see live refresh for recent tasks or alarms for any of the vCenter Server systems in the vCenter Server Single-Sign On domain.

In this example, you have two vCenter Server instances (A and B) connected to the same vCenter Server Single-Sign On domain. With each of the vCenter Server instances, you installed a vSphere Web Client instance.

You log in to vSphere Web Client A by using `https://hostnameA/vsphere-client/`.

You log in to vSphere Web Client B by using `https://hostnameB/vsphere-client/`.

You enable live refresh of recent tasks and alarms on vSphere Web Client A, and log out from it.

You can observe the following results:

- You log in to vSphere Web Client A from `https://hostnameA/vsphere-client/`. You do not see any recent tasks or alarms in the respective Recent Tasks or Alarms portlets.
- You log in to vSphere Web Client A from `https://hostnameA:9443/vsphere-client/`. You can see live refresh of recent tasks and alarms for all the users currently performing operations on both vCenter Server systems in the vCenter Server Single-Sign On domain.
- You log in to vSphere Web Client B from `https://hostnameB/vsphere-client/`. You can see the recent tasks and alarms of only operations that you perform on vCenter Server system A or vCenter Server system B. Only after you manually refresh the vSphere Web Client B you see the latest recent tasks and alarms that result from operations performed by other users on vCenter Server system A and vCenter Server system B.

## Set an Alarm in the vSphere Web Client

You can monitor inventory objects by setting alarms on them. Setting an alarm involves selecting the type of inventory object to monitor, defining when the alarm triggers, for how long the alarm is on, and defining actions that are performed as a result of the alarm being triggered. You define alarms in the alarm definition wizard. You access the alarm definition wizard from the **Monitor** tab, under **Issues**.

When you create an alarm, you select the alarm type, the type of inventory object, and the type of activity that trigger the alarm. An activity that triggers an alarm can be any of the following:

- A specific condition or a state of the inventory object.

- An event that occurs on the object.

Depending on the type of activity that you choose to monitor, the options on the Triggers page that follow the General page in the alarm definitions wizard, change.

After defining the triggers, define the actions that the trigger causes.

### Prerequisites

Log in to the vSphere Web Client.

Required Privilege: **Alarms.Create alarm** or **Alarms.Modify alarm**

- [Create or Edit Alarms](#)

To monitor your environment, you can create and modify alarm definitions in the vSphere Web Client. You can view alarm settings from any object, but you can modify settings only through the object on which the alarm is defined.

- [Specify Alarm Name, Description, and Type](#)

General settings of an alarm definition include alarm name, description, and type. You can also enable and disable the alarm from the general settings page.

- [Specify How a Condition-Based or State-Based Alarm is Triggered](#)

You can select and configure the events, states, or conditions that trigger the alarm from the Triggers page of the alarm definition wizard.

- [Specify How an Event-Based Alarm is Triggered](#)

You can specify the events, states, or conditions that trigger the alarm on the Triggers page of the alarm definition wizard. On the General page of the alarm definition wizard, if you selected a trigger for an alarm to be a specific event occurring on the inventory object, perform the following procedure.

- [Specify Alarm Actions](#)

You can define actions that the system performs when the alarm is triggered or changes status. You can enable and disable alarms and alarm actions independently of each other.

## Create or Edit Alarms

To monitor your environment, you can create and modify alarm definitions in the vSphere Web Client. You can view alarm settings from any object, but you can modify settings only through the object on which the alarm is defined.

You can access alarm definitions in the **Monitor** tab. You can create alarm definitions from the **Monitor** tab or from the object pop-up menu.

### Prerequisites

Required Privilege: **Alarms.Create alarm** or **Alarms.Modify alarm**

## Procedure

- Create or edit alarms in the **Monitor** tab.
  - a Select an inventory object, click the **Monitor** tab, and click **Issues**.
  - b Click **Alarm Definitions**.
  - c Right-click the list of alarms, and select to add or edit an alarm.
 

You cannot edit vCenter Server predefined alarms.
- Add an alarm to an object in the object navigator.
  - a Right-click an inventory object and select **Alarms > New Alarm Definition**.

## Specify Alarm Name, Description, and Type

General settings of an alarm definition include alarm name, description, and type. You can also enable and disable the alarm from the general settings page.

### Prerequisites

- Required Privilege: **Alarms.Create alarm** or **Alarms.Modify alarm**
- In the alarm definition wizard, click the General page. See [Create or Edit Alarms](#).

### Procedure

- 1 Type a name and description.
- 2 Select the type of inventory object that this alarm monitors.
- 3 Select the type of activity that this alarm monitors.

---

**Note** Depending on the type of activity that you choose to monitor, the options on the Triggers page that follow the General page in the alarm definitions wizard, change.

---

- 4 Click the **Enable this alarm** check-box if you want to enable the alarm immediately after creation.
- 5 Click **Next**.

### What to do next

Set alarm triggers.

## Specify How a Condition-Based or State-Based Alarm is Triggered

You can select and configure the events, states, or conditions that trigger the alarm from the Triggers page of the alarm definition wizard.

The option that you choose on the General page of the alarm definition wizard determine the options available on the Triggers page. An alarm definition must contain at least one trigger before you can save it.

For information about defining triggers for an event-based alarm, see [Specify How an Event-Based Alarm is Triggered](#).

You can add multiple triggers and choose whether to trigger the alarm when one or all of them become active.

#### Prerequisites

- Required Privilege: **Alarms.Create alarm** or **Alarms.Modify alarm**

#### Procedure

- 1 Select the trigger that you want to change, or click the **Add** icon to add a trigger.
- 2 Click in the **Trigger** column, and select an option from the drop-down menu.
- 3 Click in the **Operator** column, and select an option from the drop-down menu.
- 4 Click in the **Warning Condition** column, and select an option from the drop-down menu to set the threshold for triggering a warning.
- 5 Click in the **Critical Condition** column, and select an option from the drop-down menu.
- 6 Click **Next**.

You selected and configured alarm triggers.

#### What to do next

Configure actions that follow after the alarm is triggered.

## Specify How an Event-Based Alarm is Triggered

You can specify the events, states, or conditions that trigger the alarm on the Triggers page of the alarm definition wizard. On the General page of the alarm definition wizard, if you selected a trigger for an alarm to be a specific event occurring on the inventory object, perform the following procedure.

The option that you choose on the General page of the alarm definition wizard determine the options available on the Triggers page. An alarm definition must contain at least one trigger before you can save it.

For information about defining triggers for a condition-based alarm, see [Specify How a Condition-Based or State-Based Alarm is Triggered](#).

You can add multiple triggers and choose whether to trigger the alarm when one or all of them become active.

#### Prerequisites

- Required Privilege: **Alarms.Create alarm** or **Alarms.Modify alarm**

#### Procedure

- 1 Select the trigger that you want to change, or click the **Add** icon to add a trigger.
- 2 Click in the **Event** column, and select an option from the drop-down menu.

- 3 Click in the **Status** column, and select an option from the drop-down menu.
- 4 (Optional) Configure additional conditions to be met before the alarm triggers.
  - a Click the **Add** icon to add an argument.
  - b Click in the **Argument** column, and select an option from the drop-down menu.
  - c Click in the **Operator** column, and select an option from the drop-down menu.
  - d Click in the **Value** column, and enter a value into the text box.

You can add more than one argument.

- 5 Click **Next**.

You selected and configured alarm triggers.

#### What to do next

Configure actions that follow after the alarm is triggered.

## Specify Alarm Actions

You can define actions that the system performs when the alarm is triggered or changes status. You can enable and disable alarms and alarm actions independently of each other.

vCenter Server can perform alarm actions in response to triggered alarms.

#### Prerequisites

Verify that you have navigated to the Actions page of the alarm definition wizard. See [Create or Edit Alarms](#).

Required Privilege: **Alarms.Create alarm** or **Alarms.Modify alarm**

- [Send Email as an Alarm Action](#)

You can use the SMTP agent included with vCenter Server to send email notifications when alarms are triggered.

- [Send SNMP Traps as an Alarm](#)

The SNMP agent included with vCenter Server can be used to send traps when alarms are triggered on a vCenter Server instance. The default hardware health alarms send SNMP traps by default.

- [Run a Script or a Command as an Alarm Action](#)

You can configure an alarm to run a script or a command in the vSphere Web Client when the alarm is triggered.

#### Procedure

- 1 Select the action that you want to change, or click the **Add** icon to add one.
- 2 Click in the **Action** column, and select an option from the drop-down menu.

- Click in the **Configuration** column, and enter configuration information for actions that require additional information:

Option	Action
Send a notification email	Type email addresses, separated by a comma.
Migrate VM	Complete the virtual machine migration wizard.
Run a command	<p>Take one of the following actions and press <b>Enter</b>:</p> <ul style="list-style-type: none"> <li>If the command is a .exe file, enter the full path name of the command and include any parameters. For example, to run the cmd.exe command in the C:\tools directory, with the alarmName and targetName parameters, type: <b>c:\tools\cmd.exe alarmName targetName</b></li> <li>If the command is a .bat file, enter the full path name of the command as an argument to the c:\windows\system32\cmd.exe command. Include any parameters. For example, to run the cmd.bat command in the C:\tools directory, with the alarmName and targetName parameters, type: <b>c:\windows\system32\cmd.exe /c c:\tools\cmd.bat alarmName targetName</b></li> </ul> <p>For .bat files, the command and its parameters must be formatted into one string.</p>

- (Optional) For each alarm status change column, select whether the alarm should be triggered when the alarm status changes.

Some actions do not support re-triggering when alarm status change.

- For repeat actions, select the time interval for the repetition.
- Click **Finish**.

You configured the alarm general settings, triggers, and actions. The alarm monitors the object on which it is defined, as well as child objects.

## Send Email as an Alarm Action

You can use the SMTP agent included with vCenter Server to send email notifications when alarms are triggered.

### Prerequisites

Ensure that the vCenter Server SMTP agent is properly configured to send email notifications.

Required Privilege: **Alarms.Create alarm** or **Alarms.Modify alarm**

### Procedure

- On the Actions page of the alarm definition wizard, click **Add** to add an action.
- In the **Actions** column, select **Send a notification email** from the drop-down menu.
- In the **Configuration** column, enter recipient addresses. Use commas to separate multiple addresses.
- (Optional) Configure alarm transitions and frequency.

## vCenter Server Email Agent Notifications

The following tables describe the information that is included in Alarm-based and Event-based email notifications. The first table described the information included in all email notifications. The second table describes additional information that is included in Event-based notifications.

**Table 5-1. Basic SNMP Email Notification Details**

Email Entry	Description
Target	Object for which the alarm was triggered.
Old Status	Previous alarm status. Applies only to state triggers.
New Status	Current alarm status. Applies only to state triggers.
Metric Value	Threshold value that triggered the alarm. Applies only to metric condition triggers.
Alarm Definition	Alarm definition in vCenter Server, including the alarm name and status.
Description	Localized string containing a summary of the alarm. For example: Alarm New_Alarm on host1.vmware.com changed from Gray to Red.

**Table 5-2. Additional Notification Details for Alarms Triggered by Events**

Detail	Description
Event Details	VMODL event type name.
Summary	Alarm summary, including the event type, alarm name, and target object.
Date	Time and date the alarm was triggered.
UserName	Person who initiated the action that caused the event to be created. Events caused by an internal system activity do not have a UserName value.
Host	Host on which the alarm was triggered.
Resource Pool	Resource pool on which the alarm was triggered.
Datacenter	Data center on which the alarm was triggered.
Arguments	Arguments passed with the alarm and their values.

## Send SNMP Traps as an Alarm

The SNMP agent included with vCenter Server can be used to send traps when alarms are triggered on a vCenter Server instance. The default hardware health alarms send SNMP traps by default.

### Prerequisites

Ensure that vCenter Server SNMP agents and ESXi SNMP agents are properly configured.

Ensure that SNMP trap receiver agents are properly configured.

Required Privilege: **Alarms.Create alarm** or **Alarms.Modify alarm**

### Procedure

- 1 On the Actions page of the alarm definition wizard, click **Add**.
- 2 In the **Actions** column, select **Send a notification trap** from the drop-down menu.



- 3 (Optional) Configure alarm transitions and frequency.
- 4 Click **Finish** to save the alarm settings.

### SNMP Trap Notifications

The following table describes the information that is included in vCenter Server and ESXi trap notifications.

**Table 5-3. SNMP Trap Notification Details**

Trap Entry	Description
Type	The state vCenter Server is monitoring for the alarm. Options include Host Processor (or CPU) usage, Host Memory usage, Host State, Virtual Machine Processor (or CPU) usage, Virtual Machine Memory usage, Virtual Machine State, Virtual Machine Heartbeat.
Name	The name of the host or virtual machine that triggers the alarm.
Old Status	The alarm status before the alarm was triggered.
New Status	The alarm status when the alarm is triggered.
Object Value	The object value when the alarm is triggered.

### Run a Script or a Command as an Alarm Action

You can configure an alarm to run a script or a command in the vSphere Web Client when the alarm is triggered.

Use the alarm environment variables to define complex scripts and attach them to multiple alarms or inventory objects. For example, you can write a script that enters the following trouble ticket information into an external system when an alarm is triggered:

- Alarm name
- Object on which the alarm was triggered
- Event that triggered the alarm
- Alarm trigger values

When you write the script, include the following environment variables in the script:

- VMWARE\_ALARM\_NAME
- VMWARE\_ALARM\_TARGET\_NAME
- VMWARE\_ALARM\_EVENTDESCRIPTION
- VMWARE\_ALARM\_ALARMVALUE

You can attach the script to any alarm on any object without changing the script.

The script runs on the vCenter Server machine, and it runs even if you close the vSphere Web Client.

#### Prerequisites

Required Privilege: **Alarms.Create alarm** or **Alarms.Modify alarm**

## Procedure

- 1 On the Actions tab of the alarm definitions wizard, click **Add** to add an action.
- 2 In the **Actions** column, select **Run a command** from the drop-down menu.
- 3 In the **Configuration** column, type script or command information:

For this type of command...	Enter this...
<b>EXE executable files</b>	Full pathname of the command. For example, to run the <code>cmd.exe</code> command in the <code>C:\tools</code> directory, type: <code>c:\tools\cmd.exe.</code>
<b>BAT batch file</b>	Full pathname of the command as an argument to the <code>c:\windows\system32\cmd.exe</code> command. For example, to run the <code>cmd.bat</code> command in the <code>C:\tools</code> directory, type: <code>c:\windows\system32\cmd.exe /c c:\tools\cmd.bat.</code>

**Note** The command and its parameters must be formatted into one string.

If your script does not make use of the alarm environment variables, include any necessary parameters in the configuration field. Enclose parameters in curly brackets. For example:

```
c:\tools\cmd.exe {alarmName} {targetName}
c:\windows\system32\cmd.exe /c c:\tools\cmd.bat {alarmName} {targetName}
```

The script can run on any platform. You must provide the path to the script and argument keys. For example:

```
/var/myscripts/myAlarmActionScript {alarmName} {targetName}
```

- 4 (Optional) Configure alarm transitions and frequency.
- 5 Click **Finish** to save the alarm settings.

## Alarm Environment Variables for Scripts

To simplify script configuration for alarm actions, VMware provides environment variables for VMware alarms. Use the variables to define more complex scripts and attach them to multiple alarms or inventory objects so that the alarm action occurs when the alarm triggers.

**Table 5-4. Alarm Environment Variables**

Variable Name	Variable Description	Supported Alarm Type
VMWARE_ALARM_NAME	The name of the triggered alarm.	Condition, State, Event
VMWARE_ALARM_ID	The MOID of the triggered alarm.	Condition, State, Event
VMWARE_ALARM_TARGET_NAME	The name of the entity on which the alarm triggered.	Condition, State, Event
VMWARE_ALARM_TARGET_ID	The MOID of the entity on which the alarm triggered.	Condition, State, Event

**Table 5-4. Alarm Environment Variables (Continued)**

Variable Name	Variable Description	Supported Alarm Type
VMWARE_ALARM_OLDSTATUS	The old status of the alarm.	Condition, State, Event
VMWARE_ALARM_NEWSTATUS	The new status of the alarm.	Condition, State, Event
VMWARE_ALARM_TRIGGERINGSUMMARY	A multiline summary of the alarm.	Condition, State, Event
VMWARE_ALARM_DECLARINGSUMMARY	A single-line declaration of the alarm expression.	Condition, State, Event
VMWARE_ALARM_ALARMVALUE	The value that triggered the alarm.	Condition, State
VMWARE_ALARM_EVENTDESCRIPTION	A description of the alarm status change event.	Condition, State
VMWARE_ALARM_EVENTDESCRIPTION	A description of the event that triggered the alarm.	Event
VMWARE_ALARM_EVENT_USERNAME	The user name associated with the event.	Event
VMWARE_ALARM_EVENT_DATACENTER	The name of the data center in which the event occurred.	Event
VMWARE_ALARM_EVENT_COMPUTERESOURCE	The name of the cluster or resource pool in which the event occurred.	Event
VMWARE_ALARM_EVENT_HOST	The name of the host on which the event occurred.	Event
VMWARE_ALARM_EVENT_VM	The name of the virtual machine on which the event occurred.	Event
VMWARE_ALARM_EVENT_NETWORK	The name of the network on which the event occurred.	Event
VMWARE_ALARM_EVENT_DATASTORE	The name of the datastore on which the event occurred.	Event
VMWARE_ALARM_EVENT_DVS	The name of the vSphere Distributed Switch on which the event occurred.	Event

### Alarm Command-Line Parameters

VMware provides command-line parameters that function as a substitute for the default alarm environment variables. You can use these parameters when running a script as an alarm action for a condition, state, or event alarm.

The command-line parameters enable you to pass alarm information without having to change an alarm script. For example, you can use these parameters when you have an external program for which you do not have the source. You can pass in the necessary data by using the substitution parameters, which take precedence over the environment variables. You pass the parameters through the **Configuration** dialog box in the alarm definition wizard or on a command line.

**Table 5-5. Command-Line Parameters for Alarm Action Scripts**

Variable	Description
{eventDescription}	The text of the alarmStatusChange event. The {eventDescription} variable is supported only for Condition and State alarms.
{targetName}	The name of the entity on which the alarm is triggered.
{alarmName}	The name of the alarm that is triggered.
{triggeringSummary}	A summary of the alarm trigger values.
{declaringSummary}	A summary of the alarm declaration values.
{oldStatus}	The alarm status before the alarm is triggered.
{newStatus}	The alarm status after the alarm is triggered.
{target}	The inventory object on which the alarm is set.

## Set an Alarm in the vSphere Client

In the vSphere Client you define alarms in the alarm definition wizard. You access the alarm definition wizard from the **Configure** tab, under **More**.

### Create an Alarm

To monitor your environment, you can create alarm definitions in the vSphere Client . You can access alarm definitions in the **Configure** tab.

Create alarms in the **Configure** tab.

#### Prerequisites

Required Privilege: **Alarms.Create alarm** or **Alarms.Modify alarm**

#### Procedure

- 1 Select an inventory object, click the **Configure** tab, and click **More**.
- 2 Click **Alarm Definitions**.
- 3 Click **Add**.  
You cannot edit alarms in the vSphere Client.
- 4 Select an alarm, click **Enable** to enable an alarm.
- 5 Select an alarm, click **Disable** to disable an alarm.
- 6 Select an alarm, click **Delete** to delete an alarm.

### Specify Alarm Name, Description, and Target

Settings of an alarm definition include alarm name, description, and target.

**Prerequisites**

- Required Privilege: **Alarms.Create alarm** or **Alarms.Modify alarm**
- In the Alarms Definitions page, click Add. See [Create an Alarm](#)

**Procedure**

- 1 Type a name and description.
- 2 Select the type of inventory object that this alarm monitors from the **Target type** drop-down menu  
Depending on the type of target that you choose to monitor, the summary that follows the **Target**, change.
- 3 Click **Next**.

---

**Note** Depending on the type of activity that you choose to monitor, the options on the Alarm Rule page, change.

---

Set alarm rule.

**Specify How a Condition-Based or State-Based Alarm is Triggered**

You can select and configure the events, states, or conditions that trigger the alarm from the **Alarm Rule** page in the **New Alarm Definition** wizard.

An alarm definition must contain at least one trigger before you can save it.

**Prerequisites**

Required Privilege: **Alarms.Create alarm** or **Alarms.Modify alarm**

**Procedure**

- 1 Select a trigger from the drop-down menu.  
The combined event triggers are displayed. You can set the condition based on a single event only. You must create multiple alarms for multiple events.
- 2 Click **Add Argument** to select an argument from the drop-down menu.  
It supports **ALL** based expression, option to select **ANY** is not available. You must create a separate alarm definition for each trigger. The **OR** operator is not supported in the vSphere Client. However, you can combine more than one condition trigger with **AND** operator.
- 3 Select an operator from the drop-down menu.
- 4 Select an option from the drop-down menu to set the threshold for triggering an alarm
- 5 Select severity of the alarm from the drop-down menu.  
You can set the condition to either **Warning** or **Critical**, but not for both. You must create a separate alarm definition for warning and critical status.

**6** Send email notifications

- a To send email notifications when alarms are triggered, enable **Send email notifications**.
- b In the **Email to**, enter recipient addresses. Use commas to separate multiple addresses.

**7** To send traps when alarms are triggered on a vCenter Server instance, enable **Send SNMP traps**.**8** Run scripts

- a To run scripts when alarms are triggered, enable **Run script**.
- b In **Run this script** column, type script or command information:

Option	Action
Send a notification email	Type email addresses, separated by a comma.
Migrate VM	Complete the virtual machine migration wizard.
Run a command	<p>Take one of the following actions and press <b>Enter</b>:</p> <ul style="list-style-type: none"> <li>■ If the command is a .exe file, enter the full path name of the command and include any parameters. For example, to run the cmd.exe command in the C:\tools directory, with the alarmName and targetName parameters, type: <b>c:\tools\cmd.exe alarmName targetName</b></li> <li>■ If the command is a .bat file, enter the full path name of the command as an argument to the c:\windows\system32\cmd.exe command. Include any parameters. For example, to run the cmd.bat command in the C:\tools directory, with the alarmName and targetName parameters, type: <b>c:\windows\system32\cmd.exe /c c:\tools\cmd.bat alarmName targetName</b></li> </ul> <p>For .bat files, the command and its parameters must be formatted into one string.</p>

If your script does not make use of the alarm environment variables, include any necessary parameters in the configuration field. Enclose parameters in curly brackets. For example:

```
c:\tools\cmd.exe {alarmName} {targetName}
c:\windows\system32\cmd.exe /c c:\tools\cmd.bat {alarmName} {targetName}
```

The script can run on any platform. You must provide the path to the script and argument keys. For example:

```
/var/myscripts/myAlarmActionScript {alarmName} {targetName}
```

**9** (Optional) Configure alarm transitions and frequency.**What to do next**

Click **Next** to review the alarm details.

**Review and Enable Alarm**

You can review and enable the alarm in the vSphere Client

After setting the alarm rule, review the alarm before enabling the alarm.

### Prerequisites

Required Privilege: **Alarms.Create alarm** or **Alarms.Modify alarm**

### Procedure

- 1 Review the **Alarm Name**, **Description**, **Targets**, and **Alarm Rule**.
- 2 Select **Enable this alarm** to enable the alarm.

The alarm is enabled.

## Acknowledge Triggered Alarms

After you acknowledge an alarm in the vSphere Client, its alarm actions are discontinued. Alarms are not cleared, or reset when acknowledged.

Acknowledging an alarm lets other users know that you are taking ownership of the issue. For example, a host has an alarm set to monitor CPU usage. It sends an email to an administrator when the alarm is triggered. The host CPU usage spikes, triggering the alarm which sends an email to the host's administrator. The administrator acknowledges the triggered alarm to let other administrators know the problem is being addressed, and to prevent the alarm from sending more email messages. The alarm, however, is still visible in the system.

### Prerequisites

Required privilege: **Alarm.Alarm Acknowledge**

### Procedure

- Right-click the alarm in the Alarms sidebar panel and select **Acknowledge**.
- Acknowledge the alarm in the **Monitor** tab.
  - a Select an inventory object in the object navigator.
  - b Click the **Monitor** tab.
  - c Click **Issues and Alarms**, and click **Triggered Alarms**.
  - d Select an alarm and select **Acknowledge**.

## Reset Triggered Event Alarms

An alarm triggered by an event might not reset to a normal state if vCenter Server does not retrieve the event that identifies the normal condition. In such cases, reset the alarm manually in the vSphere Client to return it to a normal state.

### Prerequisites

Required privilege: **Alarm.Set Alarm Status**

## Procedure

- Right-click an alarm in the Alarms sidebar pane and select **Reset to green**.
- Reset triggered alarms in the **Monitor** tab.
  - a Select an inventory object.
  - b Click the **Monitor** tab.
  - c Click **Issues and Alarms**, and click **Triggered Alarms**.
  - d Select the alarms you want to reset.
 

Use Shift+left-click or Ctrl+left-click to select multiple alarms is supported in the vSphere Web Client.
  - e Right-click an alarm and select **Reset to Green**.

## Preconfigured vSphere Alarms

vCenter Server provides a list of default alarms, which monitor the operations of vSphere inventory objects. You must only set up actions for these alarms.

Some alarms are stateless. vCenter Server does not keep data on stateless alarms, does not compute, or display their status. Stateless alarms cannot be acknowledged or reset. Stateless alarms are indicated by an asterisk next to their name.

**Table 5-6. Default vSphere Alarms**

Alarm Name	Description
Host connection and power state	Monitors the power state of the host and whether the host is reachable.
Host CPU usage	Monitors host CPU usage.
Host memory usage	Monitors host memory usage.
Virtual machine CPU usage	Monitors virtual machine CPU usage.
Virtual machine memory usage	Monitors virtual machine memory usage.
Datastore usage on disk	Monitors datastore disk usage. <p><b>Note</b> This alarm controls the Status value for datastores in vSphere Web Client. If you disable this alarm, the datastore status is displayed as Unknown.</p>
Virtual machine CPU ready	Monitors virtual machine CPU ready time.
Virtual machine total disk latency	Monitors virtual machine total disk latency.
Virtual machine disk commands canceled	Monitors the number of virtual machine disk commands that are canceled.
Virtual machine disk reset	Monitors the number of virtual machine bus resets.
License inventory monitoring	Monitors the license inventory for compliance.
License user threshold monitoring	Monitors whether a user-defined license threshold is exceeded.



**Table 5-6. Default vSphere Alarms (Continued)**

<b>Alarm Name</b>	<b>Description</b>
License capacity monitoring	Monitors whether a license capacity is exceeded.
The host license edition is not compatible with the vCenter Server license edition	Monitors the compatibility of the vCenter Server and host license editions.
Host flash capacity exceeds the licensed limit for vSAN	Monitors whether the flash disk capacity on the host exceeds the limit of the vSAN license.
Expired vSAN license	Monitors the expiry of the vSAN license and the end of the evaluation period.
Errors occurred on the disk(s) of a vSAN host	Default alarm that monitors whether the host disks in the vSAN cluster have errors.
Timed out starting Secondary VM *	Monitors whether starting a secondary virtual machine has timed out.
No compatible host for Secondary VM	Monitors the availability of compatible hosts on which a secondary virtual machine can be created and run.
Virtual machine Fault Tolerance state changed	Monitors changes in the Fault Tolerance state of a virtual machine.
Virtual Machine Fault Tolerance vLockStep interval Status Changed	Monitors changes in the Fault Tolerance Secondary vLockStep interval.
Host processor status	Monitors the host processors.
Host memory status	Monitors host memory usage.
Host hardware fan status	Monitors host fans.
Host hardware voltage	Monitors host hardware voltage.
Host hardware temperature status	Monitors the temperature status of the host hardware.
Host hardware power status	Monitors the host power status.
Host hardware system board status	Monitors the status of host system boards.
Host battery status	Monitors the battery status of hosts.
Status of other host hardware objects	Monitors other host hardware objects.
Host storage status	Monitors host connectivity to storage devices.
Host IPMI System Event Log status	Monitors the capacity of the IPMI system event log.
Host Baseboard Management Controller status	Monitors the status of the Baseboard Management Controller.
Host error *	Monitors host error and warning events.
Virtual machine error *	Monitors virtual machine error and warning events.
Host connection failure *	Monitors host connection failures.
Unmanaged workload detected on SIOC-enabled datastore	Monitors the unmanaged I/O workload on a SIOC-enabled datastore.
Thin-provisioned volume capacity threshold exceeded	Monitors whether the thin provisioning threshold on the storage array exceeds for volumes backing the datastore.
Datastore capability alarm	Monitors the change in the capability status for volumes backing the datastore.

**Table 5-6. Default vSphere Alarms (Continued)**

<b>Alarm Name</b>	<b>Description</b>
VASA provider disconnected	Monitors the changes in the connection state of VASA providers.
VASA Provider certificate expiration alarm	Monitors whether VASA provider certificates are getting close to their expiry date.
VM storage compliance alarm	Monitors the virtual disk compliance with the object-based storage.
Datastore compliance alarm	Monitors whether the virtual disk on the datastore complies with the object-based storage.
Refreshing CA certificates and CRLs for a VASA provider failed	Monitors whether the refreshing of CA certificates and CRLs for some of the VASA providers has failed.
Insufficient vSphere HA failover resources	Monitors the sufficiency of failover cluster resources required for vSphere High Availability.
vSphere HA failover in progress	Monitors the failover progress of vSphere High Availability.
Cannot find vSphere HA master agent	Monitors whether vCenter Server is able to connect to a vSphere High Availability master agent.
vSphere HA host status	Monitors the host health status reported by vSphere High Availability.
vSphere HA virtual machine failover failed	Monitors whether a failover operation that uses vSphere High Availability failed.
vSphere HA virtual machine monitoring action	Monitors whether vSphere High Availability has restarted a virtual machine.
vSphere HA virtual machine monitoring error	Monitors whether vSphere High Availability failed to reset a virtual machine.
vSphere HA VM Component Protection could not power off a virtual machine	Monitors whether vSphere High Availability VM Component Protection cannot power off a virtual machine with an inaccessible datastore.
License error *	Monitors license errors.
Health status changed *	Monitors changes to service and extension health status.
Virtual machine component protection restart error	Monitors whether the vSphere HA VM Component Protection fails to restart a virtual machine.
Storage DRS recommendation	Monitors Storage DRS recommendations.
Storage DRS is not supported on a host	Monitors and alerts when Storage DRS is not supported on a host.
Datastore cluster is out of space	Monitors whether a datastore cluster runs out of disk space.
Datastore is in multiple datacenters	Monitors whether a datastore in a datastore cluster is visible in more than one data center.
vSphere Distributed Switch VLAN trunked status	Monitors changes in vSphere Distributed Switch VLAN trunked status.
vSphere Distributed Switch MTU matched status	Monitors changes in vSphere Distributed Switch MTU matched status.

**Table 5-6. Default vSphere Alarms (Continued)**

<b>Alarm Name</b>	<b>Description</b>
vSphere Distributed Switch MTU supported status	Monitors changes in vSphere Distributed Switch MTU supported status.
vSphere Distributed Switch teaming matched status	Monitors changes in vSphere Distributed Switch teaming matched status.
Virtual Machine network adapter reservation status	Monitors changes in the reservation status of a virtual machine network adapter.
Virtual machine Consolidation Needed status	Monitors changes in the virtual machine Consolidation Needed status.
Host virtual flash resource status	Monitors the Flash Read Cache resource status on the host.
Host virtual flash resource usage	Monitors the Flash Read Cache resource usage on the host.
Registration/unregistration of a VASA vendor provider on a vSAN host fails	Default alarm that monitors whether the registration or unregistration of a VASA vendor provider on a vSAN host fails.
Registration/unregistration of third-party IO filter storage providers fails on a host	Default alarm that monitors whether vCenter Server fails to register or unregister third-party IO filter storage providers on a host.
Service Control Agent Health Alarm	Monitors the health status of the VMware Service Control Agent.
Identity Health Alarm	Monitors the health status of the Identity Management Service.
vSphere Web Client Health Alarm	Monitors the health status of the vSphere Web Client.
ESX Agent Manager Health Alarm	Monitors the health status of the ESX Agent Manager.
Message Bus Config Health Alarm	Monitors the health status of the Message Bus Configuration Service.
Cis License Health Alarm	Monitors the health status of the License Service.
Appliance Management Health Alarm	Monitors the health status of the Appliance Management Service.
Inventory Health Alarm	Monitors the health status of the Inventory Service.
vCenter Server Health Alarm	Monitors the health status of vCenter Server.
Database Health Alarm	Monitors the database health status. When database space reaches 80%, vCenter Server displays a warning event. When database space reaches 95%, vCenter Server displays an error event and shuts down. You can clean up the database, or increase database storage capacity, and start vCenter Server. The alarm is triggered only for PostgreSQL and Microsoft SQL Server database health issues, and does not work with Oracle databases.
Data Service Health Alarm	Monitors the health status of the Data Service.
RBD Health Alarm	Monitors the health status of the vSphere Auto Deploy Waiter.
vService Manager Health Alarm	Monitors the health status of the vService Manager.

**Table 5-6. Default vSphere Alarms (Continued)**

<b>Alarm Name</b>	<b>Description</b>
Performance Charts Service Health Alarm	Monitors the health status of the Performance Charts Service.
Content Library Service Health Alarm	Monitors the health status of the VMware Content Library Service.
Transfer Service Health Alarm	Monitors the health status of the VMware Transfer Service.
VMware vSphere ESXi Dump Collector Health Alarm	Monitors the health status of the VMware vSphere ESXi Dump Collector Service.
VMware vAPI Endpoint Service Health Alarm	Monitors the health status of the VMware vAPI Endpoint Service.
VMware System and Hardware Health Manager Service Health Alarm	Monitors the health status of the VMware System and Hardware Health Manager Service.
VMware vSphere Profile-Driven Storage Service Health Alarm	Monitors the health status of the VMware vSphere Profile-Driven Storage Service.
VMware Common Logging Service Health Alarm	Monitors the health status of the VMware Common Logging Service.
VMware vFabric Postgres Service Health Alarm	Monitors the health status of the VMware vFabric Postgres Service.
ESXi Host Certificates Update Failure Status	Monitors whether the update of the ESXi host certificates failed.
ESXi Host Certificate Status	Monitors the certificate status of an ESXi host.
ESXi Host Certificate Verification Failure Status	Monitors whether the verification of an ESXi host certificate failed.
vSphere vCenter Host Certificate Management Mode	Monitors changes in the certificate management mode of vCenter Server.
Root Certificate Status	Monitors whether a root certificate is getting close to its expiration date.
GPU ECC Uncorrected Memory Alarm	Monitors the GPU ECC uncorrected memory status.
GPU ECC Corrected Memory Alarm	Monitors the GPU ECC corrected memory status.
GPU Thermal Condition Alarm	Monitors the GPU Thermal condition status.
Network connectivity lost	Monitors the network connectivity on a virtual switch.
Network uplink redundancy lost	Monitors network uplink redundancy on a virtual switch.
Network uplink redundancy degraded *	Monitors network uplink redundancy degradation on a virtual switch.
VMKernel NIC not configured correctly *	Monitors incorrectly configured VMkernel NICs.
Cannot connect to storage *	Monitors host connectivity to a storage device.
Migration error *	Monitors whether a virtual machine cannot be migrated or relocated, or is orphaned.
Exit standby error	Monitors whether a host cannot exit standby mode.

**Table 5-7. Deprecated vSphere Alarms**

<b>Alarm name</b>	<b>Description</b>
Cannot connect to network	Monitors the network connectivity on a virtual switch.
IPv6 TSO not supported	Monitors whether the IPv6 TSO packets sent by the guest operating system of a virtual machine are dropped.
SRM Consistency Group Violation	Datastore cluster has datastores that belong to different SRM consistency groups.
Virtual machine high availability error	Monitors High Availability errors on a virtual machine.
Cluster high availability error *	Monitors High Availability errors on a cluster.
Health status monitoring	Monitors changes in the overall health status of vCenter Server components.
Pre-4.1 host connected to SIOC-enabled datastore	Monitors whether a host running ESX/ESXi 4.1 or earlier is connected to a SIOC-enabled datastore.
Host service console swap rates	Monitors host service console memory swap rates.

# Monitoring Solutions with the vCenter Solutions Manager

# 6

In the vSphere Web Client, you can view an inventory of installed solutions, view detailed information about the solutions, and monitor the solution health status. A solution is an extension of vCenter Server that adds new functions to a vCenter Server instance.

VMware products that integrate with vCenter Server are also considered solutions. For example, vSphere ESX Agent Manager is a solution provided by VMware to let you manage host agents that add new capabilities to ESX/ESXi hosts.

You can install a solution to add functionality from third-party technologies to the standard functions of vCenter Server. Solutions typically are delivered as OVF packages. You can install and deploy solutions from the vSphere Web Client. You can integrate solutions into the vCenter Solutions Manager, which provides a view in the vSphere Web Client that lists all solutions.

If a virtual machine or vApp is running a solution, a custom icon represents it in the inventory of the vSphere Web Client. Each solution registers a unique icon to identify that the virtual machine or vApp is being managed by that solution. The icons show the power states (powered on, paused, or powered off). The solutions display more than one type of icon if they manage more than one type of virtual machine or vApp.

When you power on or power off a virtual machine or vApp, you are notified that you are performing this operation on an entity that is managed by the Solutions Manager. When you attempt an operation on a virtual machine or a vApp that is managed by a solution, an informational warning message appears.

For more information, see the *Developing and Deploying vSphere Solutions, vServices, and ESX Agents* documentation.

This chapter includes the following topics:

- [View Solutions and vServices](#)
- [Monitoring Agents](#)
- [Monitoring vServices](#)

## View Solutions and vServices

In the vSphere Web Client, you can view information about solutions and vService providers. A vService is a service that a solution provides to specific applications that run inside virtual machines and vApps.

## Procedure

- 1 Navigate to the vCenter Server system in the object navigator.
- 2 Double-click the vCenter Server object.
- 3 Click **Extensions**.
- 4 Select a solution.  
The **Summary** tab displays more information about the solution.
- 5 To view vService provider information, click **Monitor**, and click **vServices**.

## Monitoring Agents

The vCenter Solutions Manager displays the vSphere ESX Agent Manager agents that you use to deploy and manage related agents on ESX/ESXi hosts.

You can use the Solutions Manager to keep track of whether the agents of a solution are working as expected. Outstanding issues are reflected by the solution's ESX Agent Manager status and a list of issues.

When the status of a solution changes, the Solutions Manager updates the ESX Agent Manager summary status and state. Administrators use this status to track whether the goal state is reached.

The agent health status is indicated by a specific color.

**Table 6-1. ESX Agent Manager health status**

Status	Description
Red	The solution must intervene for the ESX Agent Manager to proceed. For example, if a virtual machine agent is powered off manually on a compute resource and the ESX Agent Manager does not attempt to power on the agent. The ESX Agent Manager reports this action to the solution, and the solution alerts the administrator to power on the agent.
Yellow	The ESX Agent Manager is actively working to reach a goal state. The goal state can be enabled, disabled, or uninstalled. For example, when a solution is registered, its status is yellow until the ESX Agent Manager deploys the solutions agents to all the specified compute resources. A solution does not need to intervene when the ESX Agent Manager reports its ESX Agent Manager health status as yellow.
Green	A solution and all its agents have reached the goal state.

## Monitoring vServices

A vService is a service or function that a solution provides to virtual machines and vApps. A solution can provide one or more vServices. These vServices integrate with the platform and are able to change the environment in which the vApp or virtual machine runs.

A vService is a type of service for a virtual machine and a vApp provided by a vCenter extension. Virtual machines and vApps can have dependencies on vServices. Each dependency is associated with a vService type. The vService type must be bound to a particular vCenter extension that implements that vService type. This vService type is similar to a virtual hardware device. For example, a virtual machine can have a networking device that at deployment must be connected to a particular network.

The vService Manager allows a solution to connect to operations related to OVF templates:

- Importing OVF templates. Receive a callback when OVF templates with a vService dependency of a certain type are imported.
- Exporting OVF templates. Inserts OVF sections when a virtual machine is exported.
- OVF environment generation. Inserts OVF sections into the OVF environment at the power-on instance.

The **vServices** tab in the Solution Manager provides details for each vCenter extension. This information allows you to monitor vService providers and list the virtual machines or vApps to which they are bound.



# Monitoring the Health of Services and Nodes

# 7

You can monitor the health status of services and nodes to determine whether problems exist in your environment.

The vSphere Web Client provides an overview of all services and nodes across the management stack of the vCenter Server system. A list of default services is available for each vCenter Server instance.

## View the Health Status of Services and Nodes

In the vSphere Web Client, you can view the health status of vCenter Server services and nodes.

vCenter Server instances and machines that run vCenter Server services are considered nodes.

Graphical badges represent the health status of services and nodes.

### Prerequisites

Verify that the user you use to log in to the vCenter Server instance is a member of the SystemConfiguration.Administrators group in the vCenter Single Sign-On domain.

### Procedure

- 1 Log in as `administrator@your_domain_name` to the vCenter Server instance by using the vSphere Web Client.
- 2 On the vSphere Web Client Home page, click **System Configuration**.

You can view the health status badges for the services and nodes.

**Table 7-1. Health States**

Badge Icon	Description
	Good. The health of the object is normal.
	Warning. The object is experiencing some problems.
	Critical. The object is either not functioning properly or will stop functioning soon.
	Unknown. No data is available for this object.

- 3 (Optional) In the Services Health and Nodes Health panes, click the hyperlink next to the health badge to view all services and nodes in this health state.

For example, in the Services Health pane, click the hyperlink of the Warning health status. In the dialog box that pops up, select a service to view more information about the service and attempt to resolve the health issues of the service.

# Performance Monitoring Utilities: resxtop and esxtop

# 8

The `resxtop` and `esxtop` command-line utilities provide a detailed look at how ESXi uses resources in real time. You can start either utility in one of three modes: interactive (default), batch, or replay.

The fundamental difference between `resxtop` and `esxtop` is that you can use `resxtop` remotely, whereas you can start `esxtop` only through the ESXi Shell of a local ESXi host.

This chapter includes the following topics:

- [Using the esxtop Utility](#)
- [Using the resxtop Utility](#)
- [Using esxtop or resxtop in Interactive Mode](#)
- [Using Batch Mode](#)
- [Using Replay Mode](#)

## Using the esxtop Utility

You can run the `esxtop` utility using the ESXi Shell to communicate with the management interface of the ESXi host. You must have root user privileges.

Type the command, using the options you want:

```
esxtop [-h] [-v] [-b] [-s] [-a] [-c config file] [-R vm-support_dir_path] [-d delay] [-n iterations]
```

The `esxtop` utility reads its default configuration from `.esxtop50rc` on the ESXi system. This configuration file consists of nine lines.

The first eight lines contain lowercase and uppercase letters to specify which fields appear in which order on the CPU, memory, storage adapter, storage device, virtual machine storage, network, interrupt, and CPU power panels. The letters correspond to the letters in the Fields or Order panels for the respective `esxtop` panel.

The ninth line contains information on the other options. Most important, if you saved a configuration in secure mode, you do not get an insecure `esxtop` without removing the `s` from the seventh line of your `.esxtop50rc` file. A number specifies the delay time between updates. As in interactive mode, typing `c`, `m`, `d`, `u`, `v`, `n`, `I`, or `p` determines the panel with which `esxtop` starts.

---

**Note** Do not edit the `.esxtop50rc` file. Instead, select the fields and the order in a running `esxtop` process, make changes, and save this file using the `W` interactive command.

---

## Using the `resxtop` Utility

The `resxtop` utility is a vSphere CLI command.

Before you can use any vSphere CLI commands, you must either download and install a vSphere CLI package or deploy the vSphere Management Assistant (vMA) to your ESXi host or vCenter Server system.

---

**Note** `resxtop` is supported only on Linux.

---

After it is set up, start `resxtop` from the command line. For remote connections, you can connect to a host either directly or through vCenter Server.

To launch `resxtop` and connect to a remote server, enter this command

```
resxtop --server <hostname> --username <user>
```

The command-line options listed in the following table are the same as for `esxtop` (except for the `R` option) with additional connection options.

---

**Note** `resxtop` does not use all the options shared by other vSphere CLI commands.

---

**Table 8-1. `resxtop` Command-Line Options**

Option	Description
[server]	Name of the remote host to connect to (required). If connecting directly to the ESXi host, use the name of that host. If your connection to the ESXi host is indirect (that is, through vCenter Server), use the name of the vCenter Server system for this option.
[vihost]	If you connect indirectly (through vCenter Server), this option should contain the name of the ESXi host you connect to. If you connect directly to the host, this option is not used. Note that the host name needs to be the same as what appears in the vSphere Web Client.
[portnumber]	Port number to connect to on the remote server. The default port is 443, and unless this is changed on the server, this option is not needed.
[username]	User name to be authenticated when connecting to the remote host. The remote server prompts you for a password.

You can also use `resxtop` on a local ESXi host by omitting the `server` option on the command line. The command defaults to `localhost`.

## Using esxtop or resxtop in Interactive Mode

By default, `resxtop` and `esxtop` run in interactive mode. Interactive mode displays statistics in different panels.

A help menu is available for each panel.

## Interactive Mode Command-Line Options

You can use various command-line options with `esxtop` and `resxtop` in interactive mode.

**Table 8-2. Interactive Mode Command-Line Options**

Option	Description
<code>h</code>	Prints help for <code>resxtop</code> (or <code>esxtop</code> ) command-line options.
<code>v</code>	Prints <code>resxtop</code> (or <code>esxtop</code> ) version number.
<code>s</code>	Calls <code>resxtop</code> (or <code>esxtop</code> ) in secure mode. In secure mode, the <code>-d</code> command, which specifies delay between updates, is disabled.
<code>d</code>	Specifies the delay between updates. The default is five seconds. The minimum is two seconds. Change this with the interactive command <code>s</code> . If you specify a delay of less than two seconds, the delay is set to two seconds.
<code>n</code>	Number of iterations. Updates the display <code>n</code> times and exits. Default value is 10000.
<code>server</code>	The name of the remote server host to connect to (required for <code>resxtop</code> only).
<code>vihost</code>	If you connect indirectly (through vCenter Server), this option should contain the name of the ESXi host you connect to. If you connect directly to the ESXi host, this option is not used. Note that the host name needs to be the same as what is displayed in the vSphere Web Client.
<code>portnumber</code>	The port number to connect to on the remote server. The default port is 443, and unless this is changed on the server, this option is not needed. ( <code>resxtop</code> only)
<code>username</code>	The user name to be authenticated when connecting to the remote host. The remote server prompts you for a password, as well ( <code>resxtop</code> only).
<code>a</code>	Show all statistics. This option overrides configuration file setups and shows all statistics. The configuration file can be the default <code>~/esxtop50rc</code> configuration file or a user-defined configuration file.
<code>c filename</code>	Load a user-defined configuration file. If the <code>-c</code> option is not used, the default configuration filename is <code>~/esxtop50rc</code> . Create your own configuration file, specifying a different filename, using the <code>W</code> single-key interactive command.

## Common Statistics Description

Several statistics appear on the different panels while `resxtop` (or `esxtop`) is running in interactive mode. These statistics are common across all four panels.

The Uptime line, found at the top of each of the four `resxtop` (or `esxtop`) panels, displays the current time, time since last reboot, number of currently running worlds and load averages. A world is an ESXi VMkernel schedulable entity, similar to a process or thread in other operating systems.

Below that the load averages over the past one, five, and fifteen minutes appear. Load averages consider both running and ready-to-run worlds. A load average of 1.00 means that there is full utilization of all physical CPUs. A load average of 2.00 means that the ESXi system might need twice as many physical CPUs as are currently available. Similarly, a load average of 0.50 means that the physical CPUs on the ESXi system are half utilized.

## Statistics Columns and Order Pages

You can define the order of fields displayed in interactive mode.

If you press `f`, `F`, `o`, or `O`, the system displays a page that specifies the field order on the top line and short descriptions of the field contents. If the letter in the field string corresponding to a field is uppercase, the field is displayed. An asterisk in front of the field description indicates whether a field is displayed.

The order of the fields corresponds to the order of the letters in the string.

From the Field Select panel, you can:

- Toggle the display of a field by pressing the corresponding letter.
- Move a field to the left by pressing the corresponding uppercase letter.
- Move a field to the right by pressing the corresponding lowercase letter.

## Interactive Mode Single-Key Commands

When running in interactive mode, `resxtop` (or `esxtop`) recognizes several single-key commands.

All interactive mode panels recognize the commands listed in the following table. The command to specify the delay between updates is disabled if the `s` option is given on the command line. All sorting interactive commands sort in descending order.

**Table 8-3. Interactive Mode Single-Key Commands**

Key	Description
<code>h</code> or <code>?</code>	Displays a help menu for the current panel, giving a brief summary of commands, and the status of secure mode.
<code>space</code>	Immediately updates the current panel.
<code>^L</code>	Erases and redraws the current panel.
<code>f</code> or <code>F</code>	Displays a panel for adding or removing statistics columns (text boxes) to or from the current panel.
<code>o</code> or <code>O</code>	Displays a panel for changing the order of statistics columns on the current panel.
<code>#</code>	Prompts you for the number of statistics rows to display. Any value greater than 0 overrides automatic determination of the number of rows to show, which is based on window size measurement. If you change this number in one <code>resxtop</code> (or <code>esxtop</code> ) panel, the change affects all four panels.
<code>s</code>	Prompts you for the delay between updates, in seconds. Fractional values are recognized down to microseconds. The default value is five seconds. The minimum value is two seconds. This command is not available in secure mode.
<code>W</code>	Write the current setup to an <code>esxtop</code> (or <code>resxtop</code> ) configuration file. This is the recommended way to write a configuration file. The default filename is the one specified by <code>-c</code> option, or <code>~/ .esxtop50rc</code> if the <code>-c</code> option is not used. You can also specify a different filename on the prompt generated by this <code>W</code> command.
<code>q</code>	Quit the interactive mode.
<code>c</code>	Switch to the CPU resource utilization panel.

**Table 8-3. Interactive Mode Single-Key Commands (Continued)**

Key	Description
p	Switch to the CPU Power utilization panel.
m	Switch to the memory resource utilization panel.
d	Switch to the storage (disk) adapter resource utilization panel.
u	Switch to storage (disk) device resource utilization screen.
v	Switch to storage (disk) virtual machine resource utilization screen.
n	Switch to the network resource utilization panel.
i	Switch to the interrupt panel.

## CPU Panel

The CPU panel displays server-wide statistics as well as statistics for the individual world, resource pool, and virtual machine CPU utilization.

Resource pools, virtual machines that are running, or other worlds are at times called groups. For worlds belonging to a virtual machine, statistics for the virtual machine that is running are displayed. All other worlds are logically aggregated into the resource pools that contain them.

**Table 8-4. CPU Panel Statistics**

Line	Description
PCPU USED(%)	<p>A PCPU refers to a physical hardware execution context. It can be a physical CPU core if the hyperthreading is unavailable or disabled, or a logical CPU (LCPU or SMT thread) if the hyperthreading is enabled.</p> <p>PCPU USED(%) displays the following percentages:</p> <ul style="list-style-type: none"> <li>■ percentage of CPU usage per PCPU</li> <li>■ percentage of CPU usage averaged over all PCPUs</li> </ul> <p>CPU Usage (%USED) is the percentage of PCPU nominal frequency that was used since the last screen update. It equals the total sum of %USED for Worlds that ran on this PCPU.</p> <hr/> <p><b>Note</b> If a PCPU is running at frequency that is higher than its nominal (rated) frequency, then PCPU USED(%) can be greater than 100%.</p> <hr/> <p>If a PCPU and its partner are busy when hyperthreading is enabled, each PCPU account for half of the CPU usage.</p>
PCPU UTIL(%)	<p>A PCPU refers to a physical hardware execution context. It can be a physical CPU core if the hyperthreading is unavailable or disabled, or a logical CPU (LCPU or SMT thread) if the hyperthreading is enabled.</p> <p>PCPU UTIL(%) represents the percentage of real time that the PCPU was not idle (raw PCPU utilization). It displays the percentage CPU utilization per PCPU, and the percentage CPU utilization averaged over all PCPUs.</p> <hr/> <p><b>Note</b> PCPU UTIL(%) might differ from PCPU USED(%) due to power management technologies or hyperthreading.</p> <hr/>
ID	Resource pool ID or virtual machine ID of the resource pool or virtual machine of the world that is running. Alternatively, the world ID of the world that is running.
GID	Resource pool ID of the resource pool or virtual machine of the world that is running.

**Table 8-4. CPU Panel Statistics (Continued)**

Line	Description
NAME	Name of the resource pool or virtual machine of the world that is running, or name of the world that is running.
NWLD	Number of members in the resource pool or virtual machine of the world that is running. If a Group is expanded using the interactive command <code>e</code> , then NWLD for all the resulting worlds is 1.
%STATE TIMES	Set of CPU statistics made up of the following percentages. For a world, the percentages are a percentage of one physical CPU core.
%USED	Percentage of physical CPU core cycles used by the resource pool, virtual machine, or world. %USED might depend on the frequency with which the CPU core is running. When running with lower CPU core frequency, %USED can be smaller than %RUN. On CPUs which support the turbo mode, CPU frequency can also be higher than the nominal (rated) frequency, and %USED can be larger than %RUN. $\%USED = \%RUN + \%SYS - \%OVRP$
%SYS	Percentage of time spent in the ESXi VMkernel on behalf of the resource pool, virtual machine, or world to process interrupts and to perform other system activities. This time is part of the time used to calculate %USED. $\%USED = \%RUN + \%SYS - \%OVRP$
%WAIT	Percentage of time the resource pool, virtual machine, or world spent in the blocked or busy wait state. This percentage includes the percentage of time the resource pool, virtual machine, or world was idle. $100\% = \%RUN + \%RDY + \%CSTP + \%WAIT$
%VMWAIT	The total percentage of time the Resource Pool/World spent in a blocked state waiting for events.
%IDLE	Percentage of time the resource pool, virtual machine, or world was idle. Subtract this percentage from %WAIT to see the percentage of time the resource pool, virtual machine, or world was waiting for some event. The difference, %WAIT - %IDLE, of the VCPU worlds can be used to estimate guest I/O wait time. To find the VCPU worlds, use the single-key command <code>e</code> to expand a virtual machine and search for the world NAME starting with "vcpu". (The VCPU worlds might wait for other events in addition to I/O events, so this measurement is only an estimate.)
%RDY	Percentage of time the resource pool, virtual machine, or world was ready to run, but was not provided CPU resources on which to execute. $100\% = \%RUN + \%RDY + \%CSTP + \%WAIT$
%MLMTD (max limited)	Percentage of time the ESXi VMkernel deliberately did not run the resource pool, virtual machine, or world because doing so would violate the resource pool, virtual machine, or world's limit setting. Because the resource pool, virtual machine, or world is ready to run when it is prevented from running in this way, the %MLMTD (max limited) time is included in %RDY time.
%SWPWT	Percentage of time a resource pool or world spends waiting for the ESXi VMkernel to swap memory. The %SWPWT (swap wait) time is included in the %WAIT time.
EVENT COUNTS/s	Set of CPU statistics made up of per second event rates. These statistics are for VMware internal use only.
CPU ALLOC	Set of CPU statistics made up of the following CPU allocation configuration parameters.
AMIN	Resource pool, virtual machine, or world attribute Reservation.
AMAX	Resource pool, virtual machine, or world attribute Limit. A value of -1 means unlimited.
ASHRS	Resource pool, virtual machine, or world attribute Shares.
SUMMARY STATS	Set of CPU statistics made up of the following CPU configuration parameters and statistics. These statistics apply only to worlds and not to virtual machines or resource pools.



**Table 8-4. CPU Panel Statistics (Continued)**

Line	Description
AFFINITY BIT MASK	Bit mask showing the current scheduling affinity for the world.
HTSHARING	Current hyperthreading configuration.
CPU	The physical or logical processor on which the world was running when resxtop (or esxtop) obtained this information.
HTQ	Indicates whether the world is quarantined or not. N means no and Y means yes.
TIMER/s	Timer rate for this world.
%OVLRLP	Percentage of system time spent during scheduling of a resource pool, virtual machine, or world on behalf of a different resource pool, virtual machine, or world while the resource pool, virtual machine, or world was scheduled. This time is not included in %SYS. For example, if virtual machine A is being scheduled and a network packet for virtual machine B is processed by the ESXi VMkernel, the time spent appears as %OVLRLP for virtual machine A and %SYS for virtual machine B. $\%USED = \%RUN + \%SYS - \%OVLRLP$
%RUN	Percentage of total time scheduled. This time does not account for hyperthreading and system time. On a hyperthreading enabled server, the %RUN can be twice as large as %USED. $\%USED = \%RUN + \%SYS - \%OVLRLP$ $100\% = \%RUN + \%RDY + \%CSTP + \%WAIT$
%CSTP	Percentage of time a resource pool spends in a ready, co-deschedule state. <b>Note</b> You might see this statistic displayed, but it is intended for VMware use only. $100\% = \%RUN + \%RDY + \%CSTP + \%WAIT$
POWER	Current CPU power consumption for a resource pool (in Watts).
%LAT_C	Percentage of time the resource pool or world was ready to run but was not scheduled to run because of CPU resource contention.
%LAT_M	Percentage of time the resource pool or world was ready to run but was not scheduled to run because of memory resource contention.
%DMD	CPU demand in percentage. It represents the average active CPU load in the past minute.
CORE UTIL(%)	Percentage of CPU cycles per core when at least one of the PCPUs in this core is unhalted, and its average over all cores. This statistic only appears when hyperthreading is enabled. In batch mode, the corresponding CORE UTIL(%) statistic is displayed for each PCPU. For example, PCPU 0 and PCPU 1 have the same the CORE UTIL(%) number, and that is the number for core 0.

You can change the display using single-key commands.

**Table 8-5. CPU Panel Single-Key Commands**

Command	Description
e	Toggles whether CPU statistics are displayed expanded or unexpanded. The expanded display includes CPU resource utilization statistics broken down by individual worlds belonging to a resource pool or virtual machine. All percentages for the individual worlds are percentage of a single physical CPU. Consider these examples: <ul style="list-style-type: none"> <li>■ If the %Used by a resource pool is 30% on a two-way server, the resource pool is utilizing 30 percent of one physical core.</li> <li>■ If the %Used by a world belonging to a resource pool is 30 percent on a two-way server, that world is utilizing 30% of one physical core.</li> </ul>
U	Sorts resource pools, virtual machines, and worlds by the resource pool's or virtual machine's %Used column. This is the default sort order.
R	Sorts resource pools, virtual machines, and worlds by the resource pool's or virtual machine's %RDY column.
N	Sorts resource pools, virtual machines, and worlds by the GID column.
V	Displays virtual machine instances only.
L	Changes the displayed length of the NAME column.

## CPU Power Panel

The CPU Power panel displays CPU Power utilization statistics.

On the CPU Power panel, statistics are arranged per physical CPU. A physical CPU is a physical hardware execution context. It is the physical CPU core when hyper-threading is unavailable or disabled, or a logical CPU (LCPU or SMT thread) when hyper-threading is enabled.

**Table 8-6. CPU Power Panel Statistics**

Line	Description
Power Usage	Current total power usage (in Watts).
Power Cap	Total power cap (in Watts).
PSTATE MHZ	Clock frequency per state.
%USED	Percentage of physical CPU nominal frequency used since the last screen update. It is the same as PCPU USED(%) shown in the CPU Screen.
%UTIL	Raw physical CPU utilization is the percentage of time that physical CPU was not idle. It is the same as PCPU UTIL(%) shown in the CPU Screen.
%Cx	Percentage of time the physical CPU spent in C-State 'x'.
%Px	Percentage of time the physical CPU spent in P-State 'x'. On systems with Processor Clocking Control, P-states are not directly visible to ESXi. The <code>esxtop</code> shows the percentage of time spent at full speed under the heading 'P0' and the percentage of time spent at any lower speed under 'P1'.
%Tx	Percentage of time the physical CPU spent in T-State 'x'.
%A/MPERF	<code>aperf</code> and <code>mperf</code> are two hardware registers used to keep track of the actual frequency and nominal frequency of the processor. Displays the real-time <code>aperf</code> to <code>mperf</code> ratio in the last <code>esxtop</code> update period. $\%A/MPERF \text{ *nominal frequency of the processor} = \text{current frequency of the processor}$

## Memory Panel

The Memory panel displays server-wide and group memory utilization statistics. As on the CPU panel, groups correspond to resource pools, running virtual machines, or other worlds that are consuming memory.

The first line, found at the top of the Memory panel displays the current time, time since last reboot, number of currently running worlds, and memory overcommitment averages. The memory overcommitment averages over the past one, five, and fifteen minutes appear. Memory overcommitment of 1.00 means a memory overcommitment of 100 percent.

**Table 8-7. Memory Panel Statistics**

Field	Description
PMEM (MB)	Displays the machine memory statistics for the server. All numbers are in megabytes.
	<b>total</b> Total amount of the machine memory in the server.
	<b>vmk</b> Amount of the machine memory being used by the ESXi VMkernel.
	<b>other</b> Amount of the machine memory being used by everything other than the ESXi VMkernel.
	<b>free</b> Amount of the machine memory that is free.
VMKMEM (MB)	Displays the machine memory statistics for the ESXi VMkernel. All numbers are in megabytes.
	<b>managed</b> Total amount of the machine memory managed by the ESXi VMkernel.
	<b>min free</b> Minimum amount of the machine memory that the ESXi VMkernel aims to keep free.
	<b>rsvd</b> Total amount of the machine memory currently reserved by resource pools.
	<b>ursvd</b> Total amount of the machine memory currently unreserved.
NUMA (MB)	Displays the ESXi NUMA statistics. This line appears only if the ESXi host is running on a NUMA server. All numbers are in megabytes.
	<p>For each NUMA node in the server, two statistics are displayed:</p> <ul style="list-style-type: none"> <li>■ The total amount of machine memory in the NUMA node that is managed by ESXi.</li> <li>■ The amount of machine memory in the node that is currently free (in parentheses).</li> </ul> <p>Shared memory for the ESXi host might be larger than the total amount of memory if memory is over-committed.</p>

**Table 8-7. Memory Panel Statistics (Continued)**

<b>Field</b>	<b>Description</b>
PSHARE (MB)	Displays the ESXi page-sharing statistics. All numbers are in megabytes.
	<b>shared</b> Amount of the physical memory that is being shared.
	<b>common</b> Amount of the machine memory that is common across worlds.
	<b>saving</b> Amount of the machine memory that is saved because of page sharing.
	shared = common + saving
SWAP (MB)	Displays the ESXi swap usage statistics. All numbers are in megabytes.
	<b>curr</b> Current swap usage.
	<b>rclmtgt</b> Where the ESXi system expects the reclaimed memory to be. Memory can be reclaimed by swapping or compression.
	<b>r/s</b> Rate at which the memory is swapped in by the ESXi system from disk.
	<b>w/s</b> Rate at which the memory is swapped to disk by the ESXi system.
ZIP (MB)	Displays the ESXi memory compression statistics. All numbers are in megabytes.
	<b>zipped</b> Total compressed physical memory.
	<b>saved</b> Saved memory by compression.
MEMCTL (MB)	Displays the memory balloon statistics. All numbers are in megabytes.
	<b>curr</b> Total amount of the physical memory reclaimed using the <code>vmmemctl</code> module.
	<b>target</b> Total amount of physical memory the ESXi host attempts to reclaim using the <code>vmmemctl</code> module.
	<b>max</b> Maximum amount of the physical memory the ESXi host can reclaim using the <code>vmmemctl</code> module.
AMIN	Memory reservation for this resource pool or virtual machine.
AMAX	Memory limit for this resource pool or virtual machine. A value of -1 means Unlimited.
ASHRS	Memory shares for this resource pool or virtual machine.
NHN	Current home node for the resource pool or virtual machine. This statistic is applicable only on NUMA systems. If the virtual machine has no home node, a dash (-) appears.
NRMEM (MB)	Current amount of remote memory allocated to the virtual machine or resource pool. This statistic is applicable only on NUMA systems.
N% L	Current percentage of memory allocated to the virtual machine or resource pool that is local.
MEMSZ (MB)	Amount of physical memory allocated to a resource pool or virtual machine. The values are the same for the VMM and VMX groups. MEMSZ = GRANT + MCTLSZ + SWCUR + "never touched"

**Table 8-7. Memory Panel Statistics (Continued)**

Field	Description
GRANT (MB)	Amount of guest physical memory mapped to a resource pool or virtual machine. The consumed host machine memory is equal to GRANT - SHRDSVD. The values are the same for the VMM and VMX groups.
CNSM	Amount of the memory currently consumed by the virtual machine. The memory currently consumed by the virtual machine is equal to the amount of memory that the VM guest operating system currently uses, excluding the amount of memory saved for sharing if memory sharing is enabled on the VM, excluding the amount of memory saved if some of the VM memory is compressed. For more information on memory sharing and memory compression, see the <i>vSphere Resource Management</i> documentation.
SZTGT (MB)	Amount of machine memory the ESXi VMkernel wants to allocate to a resource pool or virtual machine. The values are the same for the VMM and VMX groups.
TCHD (MB)	Working set estimate for the resource pool or virtual machine. The values are the same for the VMM and VMX groups.
%ACTV	Percentage of guest physical memory that is being referenced by the guest. This is an instantaneous value.
%ACTVS	Percentage of guest physical memory that is being referenced by the guest. This is a slow moving average.
%ACTVF	Percentage of guest physical memory that is being referenced by the guest. This is a fast moving average.
%ACTVN	Percentage of guest physical memory that is being referenced by the guest. This is an estimation. (You might see this statistic displayed, but it is intended for VMware use only.)
MCTL?	Memory balloon driver is installed or not. <b>N</b> means no, <b>Y</b> means yes.
MCTLSZ (MB)	Amount of physical memory reclaimed from the resource pool by way of ballooning.
MCTLTGT (MB)	Amount of physical memory the ESXi system attempts to reclaim from the resource pool or virtual machine by way of ballooning.
MCTLMAX (MB)	Maximum amount of physical memory the ESXi system can reclaim from the resource pool or virtual machine by way of ballooning. This maximum depends on the guest operating system type.
SWCUR (MB)	Current swap usage by this resource pool or virtual machine.
SWTGT (MB)	Target where the ESXi host expects the swap usage by the resource pool or virtual machine to be.
SWR/s (MB)	Rate at which the ESXi host swaps in memory from disk for the resource pool or virtual machine.
SWW/s (MB)	Rate at which the ESXi host swaps resource pool or virtual machine memory to disk.
LLSWR/s (MB)	Rate at which memory is read from the host cache. The reads and writes are attributed to the VMM group only, LLSWAP statistics are not displayed for the VM group.
LLSWW/s (MB)	Rate at which memory is written to the host cache from various sources. The reads and writes are attributed to the VMM group only, LLSWAP statistics are not displayed for the VM group.
CPTRD (MB)	Amount of data read from checkpoint file.
CPTTGT (MB)	Size of checkpoint file.
ZERO (MB)	Resource pool or virtual machine physical pages that are zeroed.
SHRD (MB)	Resource pool or virtual machine physical pages that are shared.
SHRDSVD (MB)	Machine pages that are saved because of resource pool or virtual machine shared pages.

**Table 8-7. Memory Panel Statistics (Continued)**

Field	Description
OVHD (MB)	Current space overhead for resource pool.
OVHDMAX (MB)	Maximum space overhead that might be incurred by resource pool or virtual machine.
OVHDUW (MB)	Current space overhead for a user world. (You might see this statistic displayed, but it is intended for VMware use only.)
GST_NDx (MB)	Guest memory allocated for a resource pool on NUMA node x. This statistic is applicable on NUMA systems only.
OVD_NDx (MB)	VMM overhead memory allocated for a resource pool on NUMA node x. This statistic is applicable on NUMA systems only.
TCHD_W (MB)	Write working set estimate for resource pool.
CACHESZ (MB)	Compression memory cache size.
CACHEUSD (MB)	Used compression memory cache.
ZIP/s (MB/s)	Compressed memory per second.
UNZIP/s (MB/s)	Decompressed memory per second.

**Table 8-8. Memory Panel Interactive Commands**

Command	Description
M	Sort resource pools or virtual machines by MEMSZ column. This is the default sort order.
B	Sort resource pools or virtual machines by Group Memctl column.
N	Sort resource pools or virtual machines by GID column.
V	Display virtual machine instances only.
L	Changes the displayed length of the NAME column.

## Storage Adapter Panel

Statistics in the Storage Adapter panel are aggregated per storage adapter by default. Statistics can also be viewed per storage path.

**Table 8-9. Storage Adapter Panel Statistics**

Column	Description
ADAPTR	Name of the storage adapter.
PATH	Storage path name. This name is only visible if the corresponding adapter is expanded. See interactive command <a href="#">e</a> in <a href="#">Table 8-10</a> .
NPTH	Number of paths.
AQLEN	Current queue depth of the storage adapter.
CMDS/s	Number of commands issued per second.
READS/s	Number of read commands issued per second.
WRITES/s	Number of write commands issued per second.
MBREAD/s	Megabytes read per second.

**Table 8-9. Storage Adapter Panel Statistics (Continued)**

Column	Description
MBWRTN/s	Megabytes written per second.
RESV/s	Number of SCSI reservations per second.
CONS/s	Number of SCSI reservation conflicts per second.
DAVG/cmd	Average device latency per command, in milliseconds.
KAVG/cmd	Average ESXi VMkernel latency per command, in milliseconds.
GAVG/cmd	Average virtual machine operating system latency per command, in milliseconds.
QAVG/cmd	Average queue latency per command, in milliseconds.
DAVG/rd	Average device read latency per read operation, in milliseconds.
KAVG/rd	Average ESXi VMkernel read latency per read operation, in milliseconds.
GAVG/rd	Average guest operating system read latency per read operation, in milliseconds.
QAVG/rd	Average queue latency per read operation, in milliseconds.
DAVG/wr	Average device write latency per write operation, in milliseconds.
KAVG/wr	Average ESXi VMkernel write latency per write operation, in milliseconds.
GAVG/wr	Average guest operating system write latency per write operation, in milliseconds.
QAVG/wr	Average queue latency per write operation, in milliseconds.
FCMDS/s	Number of failed commands issued per second.
FREAD/s	Number of failed read commands issued per second.
FWRITE/s	Number of failed write commands issued per second.
FMBRD/s	Megabytes of failed read operations per second.
FMBWR/s	Megabytes of failed write operations per second.
FRESV/s	Number of failed SCSI reservations per second.
ABRTS/s	Number of commands aborted per second.
RESETS/s	Number of commands reset per second.
PAECMD/s	The number of PAE (Physical Address Extension) commands per second.
PAECP/s	The number of PAE copies per second.
SPLTCMD/s	The number of split commands per second.
SPLTCP/s	The number of split copies per second.

The following table displays the interactive commands you can use with the storage adapter panel.

**Table 8-10. Storage Adapter Panel Interactive Commands**

Command	Description
e	Toggles whether storage adapter statistics appear expanded or unexpanded. Allows you to view storage resource utilization statistics broken down by individual paths belonging to an expanded storage adapter. You are prompted for the adapter name.
r	Sorts by READS/s column.

**Table 8-10. Storage Adapter Panel Interactive Commands (Continued)**

Command	Description
w	Sorts by WRITES/s column.
R	Sorts by MBREAD/s read column.
T	Sorts by MBWRN/s written column.
N	Sorts first by ADAPTR column, then by PATH column. This is the default sort order.

## Storage Device Panel

The storage device panel displays server-wide storage utilization statistics.

By default, the information is grouped per storage device. You can also group the statistics per path, per world, or per partition.

**Table 8-11. Storage Device Panel Statistics**

Column	Description
DEVICE	Name of the storage device.
PATH	Path name. This name is visible only if the corresponding device is expanded to paths. See the interactive command p in <a href="#">Table 8-12</a> .
WORLD	World ID. This ID is visible only if the corresponding device is expanded to worlds. See the interactive command e in <a href="#">Table 8-12</a> . The world statistics are per world per device.
PARTITION	Partition ID. This ID is visible only if the corresponding device is expanded to partitions. See interactive command t in <a href="#">Table 8-12</a> .
NPH	Number of paths.
NWD	Number of worlds.
NPN	Number of partitions.
SHARES	Number of shares. This statistic is applicable only to worlds.
BLKSZ	Block size in bytes.
NUMBLKS	Number of blocks of the device.
DQLEN	Current device queue depth of the storage device.
WQLEN	World queue depth. This is the maximum number of ESXi VMkernel active commands that the world is allowed to have. This is a per device maximum for the world. It is valid only if the corresponding device is expanded to worlds.
ACTV	Number of commands in the ESXi VMkernel that are currently active. This statistic applies to only worlds and devices.
QUED	Number of commands in the ESXi VMkernel that are currently queued. This statistic applies to only worlds and devices.
%USD	Percentage of the queue depth used by ESXi VMkernel active commands. This statistic applies to only worlds and devices.
LOAD	Ratio of ESXi VMkernel active commands plus ESXi VMkernel queued commands to queue depth. This statistic applies to only worlds and devices.
CMDS/s	Number of commands issued per second.



**Table 8-11. Storage Device Panel Statistics (Continued)**

<b>Column</b>	<b>Description</b>
READS/s	Number of read commands issued per second.
WRITES/s	Number of write commands issued per second.
MBREAD/s	Megabytes read per second.
MBWRTN/s	Megabytes written per second.
DAVG/cmd	Average device latency per command in milliseconds.
KAVG/cmd	Average ESXi VMkernel latency per command in milliseconds.
GAVG/cmd	Average guest operating system latency per command in milliseconds.
QAVG/cmd	Average queue latency per command in milliseconds.
DAVG/rd	Average device read latency per read operation in milliseconds.
KAVG/rd	Average ESXi VMkernel read latency per read operation in milliseconds.
GAVG/rd	Average guest operating system read latency per read operation in milliseconds.
QAVG/rd	Average queue read latency per read operation in milliseconds.
DAVG/wr	Average device write latency per write operation in milliseconds.
KAVG/wr	Average ESXi VMkernel write latency per write operation in milliseconds.
GAVG/wr	Average guest operating system write latency per write operation in milliseconds.
QAVG/wr	Average queue write latency per write operation in milliseconds.
ABRTS/s	Number of commands aborted per second.
RESETS/s	Number of commands reset per second.
PAECMD/s	Number of PAE commands per second. This statistic applies to only paths.
PAECP/s	Number of PAE copies per second. This statistic applies to only paths.
SPLTCMD/s	Number of split commands per second. This statistic applies to only paths.
SPLTCP/s	Number of split copies per second. This statistic applies to only paths.

The following table displays the interactive commands you can use with the storage device panel.

**Table 8-12. Storage Device Panel Interactive Commands**

<b>Command</b>	<b>Description</b>
e	Expand or roll up storage world statistics. This command allows you to view storage resource utilization statistics separated by individual worlds belonging to an expanded storage device. You are prompted for the device name. The statistics are per world per device.
P	Expand or roll up storage path statistics. This command allows you to view storage resource utilization statistics separated by individual paths belonging to an expanded storage device. You are prompted for the device name.
t	Expand or roll up storage partition statistics. This command allows you to view storage resource utilization statistics separated by individual partitions belonging to an expanded storage device. You are prompted for the device name.
r	Sort by READS/s column.

**Table 8-12. Storage Device Panel Interactive Commands (Continued)**

Command	Description
w	Sort by WRITES/s column.
R	Sort by MBREAD/s column.
T	Sort by MBWRTN column.
N	Sort first by DEVICE column, then by PATH, WORLD, and PARTITION column. This is the default sort order.
L	Changes the displayed length of the DEVICE column.

## Virtual Machine Storage Panel

This panel displays virtual machine-centric storage statistics.

By default, statistics are aggregated on a per-resource-pool basis. One virtual machine has one corresponding resource pool, so the panel displays statistics on a per-virtual-machine basis. You can also view the statistics on per-VSCSI-device basis.

**Table 8-13. Virtual Machine Storage Panel Statistics**

Column	Description
ID	Resource pool ID or VSCSI ID of VSCSI device.
GID	Resource pool ID.
VMNAME	Name of the resource pool.
VSCSINAME	Name of the VSCSI device.
NDK	Number of VSCSI devices
CMDS/s	Number of commands issued per second.
READS/s	Number of read commands issued per second.
WRITES/s	Number of write commands issued per second.
MBREAD/s	Megabytes read per second.
MBWRTN/s	Megabytes written per second.
LAT/rd	Average latency (in milliseconds) per read.
LAT/wr	Average latency (in milliseconds) per write.

The following table lists the interactive commands you can use with the virtual machine storage panel.

**Table 8-14. Virtual Machine Storage Panel Interactive Commands**

Command	Description
e	Expand or roll up storage VSCSI statistics. Allows you to view storage resource utilization statistics broken down by individual VSCSI devices belonging to a group. You are prompted to enter the group ID. The statistics are per VSCSI device.
r	Sort by READS/s column.
w	Sort by WRITES/s column.
R	Sort by MBREAD/s column.

**Table 8-14. Virtual Machine Storage Panel Interactive Commands (Continued)**

Command	Description
T	Sort by MBWRTN/s column.
N	Sort first by VMNAME column, and then by VSCSINAME column. It is the default sort order.

## Network Panel

The Network panel displays server-wide network utilization statistics.

Statistics are arranged by port for each virtual network device configured. For physical network adapter statistics, see the row in the table that corresponds to the port to which the physical network adapter is connected. For statistics on a virtual network adapter configured in a particular virtual machine, see the row corresponding to the port to which the virtual network adapter is connected.

**Table 8-15. Network Panel Statistics**

Column	Description
PORT-ID	Virtual network device port ID.
UPLINK	Y means that the corresponding port is an uplink. N means it is not.
UP	Y means that the corresponding link is up. N means it is not.
SPEED	Link speed in Megabits per second.
FDUPLX	Y means the corresponding link is operating at full duplex. N means it is not.
USED-BY	Virtual network device port user.
DTYP	Virtual network device type. H means HUB and S means switch.
DNAME	Virtual network device name.
PKTTX/s	Number of packets transmitted per second.
PKTRX/s	Number of packets received per second.
MbTX/s	MegaBits transmitted per second.
MbRX/s	MegaBits received per second.
%DRPTX	Percentage of transmit packets dropped.
%DRPRX	Percentage of receive packets dropped.
TEAM-PNIC	Name of the physical NIC used for the team uplink.
PKTTXMUL/s	Number of multicast packets transmitted per second.
PKTRXMUL/s	Number of multicast packets received per second.
PKTTXBRD/s	Number of broadcast packets transmitted per second.
PKTRXBRD/s	Number of broadcast packets received per second.

The following table displays the interactive commands you can use with the network panel.

**Table 8-16. Network Panel Interactive Commands**

Command	Description
T	Sorts by Mb Tx column.
R	Sorts by Mb Rx column.
t	Sorts by Packets Tx column.
r	Sorts by Packets Rx column.
N	Sorts by PORT-ID column. This is the default sort order.
L	Changes the displayed length of the DNAME column.

## Interrupt Panel

The interrupt panel displays information about the use of interrupt vectors.

**Table 8-17. Interrupt Panel Statistics**

Column	Description
VECTOR	Interrupt vector ID.
COUNT/s	Total number of interrupts per second. This value is cumulative of the count for every CPU.
COUNT_x	Interrupts per second on CPU x.
TIME/int	Average processing time per interrupt (in microseconds).
TIME_x	Average processing time per interrupt on CPU x (in microseconds).
DEVICES	Devices that use the interrupt vector. If the interrupt vector is not enabled for the device, its name is enclosed in angle brackets (< and >).

## Using Batch Mode

Batch mode allows you to collect and save resource utilization statistics in a file.

After you prepare for batch mode, you can use `esxtop` or `resxtop` in this mode.

### Prepare for Batch Mode

To run in batch mode, you must first prepare for batch mode.

#### Procedure

- 1 Run `resxtop` (or `esxtop`) in interactive mode.
- 2 In each of the panels, select the columns you want.
- 3 Save this configuration to a file (by default `~/ .esxtop50rc`) using the `W` interactive command.

You can now use `resxtop` (or `esxtop`) in batch mode.

### Use `esxtop` or `resxtop` in Batch Mode

After you have prepared for batch mode, you can use `esxtop` or `resxtop` in this mode.

## Procedure

- 1 Start `resxtop` (or `esxtop`) to redirect the output to a file.

For example:

```
esxtop -b > my_file.csv
```

The filename must have a `.csv` extension. The utility does not enforce this, but the post-processing tools require it.

- 2 Process statistics collected in batch mode using tools such as Microsoft Excel and Perfmon.

In batch mode, `resxtop` (or `esxtop`) does not accept interactive commands. In batch mode, the utility runs until it produces the number of iterations requested (see command-line option `n`, below, for more details), or until you end the process by pressing `Ctrl+c`.

## Batch Mode Command-Line Options

You can use batch mode with command-line options.

**Table 8-18. Command-Line Options in Batch Mode**

Option	Description
<code>a</code>	Show all statistics. This option overrides configuration file setups and shows all statistics. The configuration file can be the default <code>~/ .esxtop50rc</code> configuration file or a user-defined configuration file.
<code>b</code>	Runs <code>resxtop</code> (or <code>esxtop</code> ) in batch mode.
<code>c filename</code>	Load a user-defined configuration file. If the <code>-c</code> option is not used, the default configuration filename is <code>~/ .esxtop41rc</code> . Create your own configuration file, specifying a different filename, using the <code>W</code> single-key interactive command.
<code>d</code>	Specifies the delay between statistics snapshots. The default is five seconds. The minimum is two seconds. If a delay of less than two seconds is specified, the delay is set to two seconds.
<code>n</code>	Number of iterations. <code>resxtop</code> (or <code>esxtop</code> ) collects and saves statistics this number of times, and then exits.
<code>server</code>	The name of the remote server host to connect to (required, <code>resxtop</code> only).
<code>vihost</code>	If you connect indirectly (through vCenter Server), this option should contain the name of the ESXi host you connect to. If you connect directly to the ESXi host, this option is not used.
<b>Note</b> The host name needs to be the same as what appears in the vSphere Web Client.	
<code>portnumber</code>	The port number to connect to on the remote server. The default port is 443, and unless this is changed on the server, this option is not needed. ( <code>resxtop</code> only)
<code>username</code>	The user name to be authenticated when connecting to the remote host. You are prompted by the remote server for a password, as well ( <code>resxtop</code> only).

## Using Replay Mode

In replay mode, `esxtop` replays resource utilization statistics collected using `vm-support`.

After you prepare for replay mode, you can use `esxtop` in this mode. See the `vm-support` man page.

In replay mode, `esxtop` accepts the same set of interactive commands as in interactive mode and runs until no more snapshots are collected by `vm-support` to be read or until the requested number of iterations are completed.

## Prepare for Replay Mode

To run in replay mode, you must prepare for replay mode.

### Procedure

- 1 Run `vm-support` in snapshot mode in the ESXi Shell.

Use the following command.

```
vm-support -S -d duration -I interval
```

- 2 Unzip and untar the resulting tar file so that `esxtop` can use it in replay mode.

You can now use `esxtop` in replay mode.

## Use esxtop in Replay Mode

You can use `esxtop` in replay mode.

Replay mode can be run to produce output in the same style as batch mode (see the command-line option `b`, below).

---

**Note** Batch output from `esxtop` cannot be played back by `resxtop`.

---

Snapshots collected by `vm-supported` can be replayed by `esxtop`. However, `vm-support` output generated by ESXi can be replayed only by `esxtop` running on the same version of ESXi.

### Procedure

- ◆ To activate replay mode, enter the following at the command-line prompt.

```
esxtop -R vm-support_dir_path
```

## Replay Mode Command-Line Options

You can use replay mode with command-line options.

The following table lists the command-line options available for `esxtop` replay mode.

**Table 8-19. Command-Line Options in Replay Mode**

Option	Description
R	Path to the vm-support collected snapshot's directory.
a	Show all statistics. This option overrides configuration file setups and shows all statistics. The configuration file can be the default <code>~/ .esxtop50rc</code> configuration file or a user-defined configuration file.
b	Runs <code>esxtop</code> in Batch mode.
c <i>filename</i>	Load a user-defined configuration file. If the <code>-c</code> option is not used, the default configuration filename is <code>~/ .esxtop50rc</code> . Create your own configuration file and specify a different filename using the <code>w</code> single-key interactive command.
d	Specifies the delay between panel updates. The default is five seconds. The minimum is two seconds. If a delay of less than two seconds is specified, the delay is set to two seconds.
n	Number of iterations <code>esxtop</code> updates the display this number of times and then exits.

# Using the vimtop Plug-In to Monitor the Resource Use of Services

# 9

You can use the `vimtop` utility plug-in to monitor vSphere services that run in the vCenter Server Appliance.

`vimtop` is a tool similar to `esxtop`, which runs in the environment of the vCenter Server Appliance. By using the text-based interface of `vimtop` in the appliance shell, you can view overall information about the vCenter Server Appliance, and a list of vSphere services and their resource use.

This chapter includes the following topics:

- [Monitor Services By Using vimtop in Interactive Mode](#)
- [Interactive Mode Command-Line Options](#)
- [Interactive Mode Single-Key Commands for vimtop](#)

## Monitor Services By Using vimtop in Interactive Mode

You can use the `vimtop` plug-in to monitor services in real time.

The default view of the `vimtop` interactive mode consists of the overview tables and the main table. You can use single-key commands in interactive mode to switch the view from processes to disks or network.

### Procedure

- 1 From an SSH client application, log in to the vCenter Server Appliance shell.
- 2 Run the `vimtop` command to access the plug-in in interactive mode.

## Interactive Mode Command-Line Options

You can use various command-line options when you run the `vimtop` command to enter the plug-in interactive mode.

**Table 9-1. Interactive Mode Command-Line Options**

Option	Description
-h	Prints help for the <code>vimtop</code> command-line options.
-v	Prints the <code>vimtop</code> version number.



**Table 9-1. Interactive Mode Command-Line Options (Continued)**

Option	Description
<code>-c filename</code>	Loads a user-defined <code>vimtop</code> configuration file. If the <code>-c</code> option is not used, the default configuration file is <code>/root/vimtop/vimtop.xml</code> . You can create your own configuration file, specifying a different filename and path by using the <code>W</code> single-key interactive command.
<code>-n number</code>	Sets the number of performed iterations before the <code>vimtop</code> exits interactive mode. <code>vimtop</code> updates the display <code>number</code> number of times and exits. The default value is 10000.
<code>-p / -d seconds</code>	Sets the update period in seconds.

## Interactive Mode Single-Key Commands for vimtop

When running in interactive mode, `vimtop` recognizes several single-key commands.

All interactive mode panels recognize the commands listed in the following table.

**Table 9-2. Interactive Mode Single-Key Commands**

Key Names	Description
<code>h</code>	Show a help menu for the current panel, giving a brief summary of commands, and the status of secure mode.
<code>i</code>	Show or hide the top line view of the overview panel of the <code>vimtop</code> plug-in.
<code>t</code>	Show or hide the Tasks section, which displays information in the overview panel about the tasks currently running on the vCenter Server instance.
<code>m</code>	Show or hide the Memory section in the overview panel.
<code>f</code>	Show or hide the CPU section which displays information in the overview panel about all available CPUs.
<code>g</code>	Show or hide the CPUs section which displays information in the overview panel about the top 4 physical CPUs.
spacebar	Immediately refreshes the current pane.
<code>p</code>	Pause the displayed information about the services resource use in the current panels.
<code>r</code>	Refresh the displayed information about the services resource use in the current panels.
<code>s</code>	Set refresh period.
<code>q</code>	Exit the interactive mode of the <code>vimtop</code> plug-in.
<code>k</code>	Displays the Disks view of the main panel.
<code>o</code>	Switch the main panel to Network view.
Esc	Clear selection or return to the Processes view of the main panel.
Enter	Select a service to view additional details.
<code>n</code>	Show or hide names of the headers in the main panel.
<code>u</code>	Show or hide the measurement units in the headers in the main panel.
left, right arrows	Select columns.
up, down arrows	Select rows.
<code>&lt;, &gt;</code>	Move a selected column.

**Table 9-2. Interactive Mode Single-Key Commands (Continued)**

---

<b>Key Names</b>	<b>Description</b>
Delete	Remove selected column.
c	Add a column to the current view of the main panel. Use spacebar to add or remove columns from the displayed list.
a	Sort the selected column in ascending order.
d	Sort the selected column in descending order.
z	Clear the sort order for all columns.
l	Set width for the selected column.
x	Return the column widths to their default values.
+	Expand selected item.
-	Collapse selected item.
w	Write the current setup to a vimtop configuration file. The default file name is the one specified by <code>-c</code> option, or <code>/root/vimtop/vimtop.xml</code> if the <code>-c</code> option is not used. You can also specify a different file name on the prompt generated by the <code>w</code> command.

# Monitoring Networked Devices with SNMP and vSphere

# 10

Simple Network Management Protocol (SNMP) is commonly used by management programs to monitor a variety of networked devices.

vSphere systems run SNMP agents, which can provide information to a management program in at least one of the following ways:

- In response to a GET, GETBULK, or GETNEXT operation, which is a specific request for information from the management system.
- By sending a notification which is an alert sent by the SNMP agent to notify the management system of a particular event or condition.

Management Information Base (MIB) files define the information that can be provided by managed devices. The MIB files define managed objects, described by object identifiers (OIDs) and variables arranged in a hierarchy.

vCenter Server and ESXi have SNMP agents. The agent provided with each product has different capabilities.

This chapter includes the following topics:

- [Using SNMP Traps with vCenter Server](#)
- [Configure SNMP for ESXi](#)
- [SNMP Diagnostics](#)
- [Monitor Guest Operating Systems with SNMP](#)
- [VMware MIB Files](#)
- [SNMPv2 Diagnostic Counters](#)

## Using SNMP Traps with vCenter Server

The SNMP agent included with vCenter Server can be used to send traps when vCenter Server starts and when an alarm is triggered on vCenter Server. The vCenter Server SNMP agent functions only as a trap emitter and does not support other SNMP operations, such as receiving GET, GETBULK, and GETNEXT requests.

vCenter Server can send SNMPv1 traps to other management applications. You must configure your management server to interpret the SNMP traps sent by vCenter Server.

To use the vCenter Server SNMP traps, configure the SNMP settings on vCenter Server and your management client software to accept the traps from vCenter Server.

The traps sent by vCenter Server are defined in `VMWARE-VC-EVENT-MIB.mib`.

## Configure SNMP Settings for vCenter Server

If you plan to use SNMP with vCenter Server, you must use the vSphere Web Client to configure the SNMP settings.

### Prerequisites

- Verify that the vSphere Web Client is connected to a vCenter Server instance.
- Verify that you have the domain name or IP address of the SNMP receiver, the port number of the receiver, and the community string.

### Procedure

- 1 In the vSphere Web Client, navigate to a vCenter Server instance.
- 2 Click the **Configure** tab.
- 3 Under Settings, click **General**.
- 4 On the vCenter Server Settings central pane, click **Edit**.

The **Edit vCenter Server Settings** wizard opens.

- 5 Click **SNMP receivers** to edit their settings.
- 6 Enter the following information for the primary receiver of the SNMP traps.

Option	Description
<b>Primary Receiver URL</b>	Enter the domain name or IP address of the receiver of SNMP traps.
<b>Enable receiver</b>	Select the check box to enable the SNMP receiver.
<b>Receiver port</b>	Enter the port number of the receiver to which the SNMP agent sends traps. If the port value is empty, vCenter Server uses port 162 by default.
<b>Community string</b>	Enter the community string that is used for authentication.

- 7 (Optional) Enter information about other SNMP receivers in the **Receiver 2 URL**, **Receiver 3 URL**, and **Receiver 4 URL** options, and select **Enabled**.
- 8 Click **OK**.

The vCenter Server system is now ready to send traps to the management system you have specified.

### What to do next

Configure your SNMP management software to receive and interpret data from the vCenter Server SNMP agent. See [Configure SNMP Management Client Software](#).

## Configure SNMP for ESXi

ESXi includes an SNMP agent that can send notifications (traps and informs) and receive GET, GETBULK, and GETNEXT requests.

In ESXi 5.1 and later releases, the SNMP agent adds support for version 3 of the SNMP protocol, offering increased security and improved functionality, including the ability to send informs. You can use `esxcli` commands to enable and configure the SNMP agent. You configure the agent differently depending on whether you want to use SNMP v1/v2c or SNMP v3.

As an alternative to configuring SNMP manually using `esxcli` commands, you can use host profiles to configure SNMP for an ESXi host. See the *vSphere Host Profiles* documentation for more information.

---

**Note** For information on configuring SNMP for ESXi 5.0 or earlier or ESX 4.1 or earlier, see the documentation for the appropriate product version.

---

- [Configure the SNMP Agent for Polling](#)

If you configure the ESXi SNMP agent for polling, it can listen for and respond to requests from SNMP management client systems, such as GET, GETNEXT and GETBULK requests.

- [Configure ESXi for SNMPv1 and SNMPv2c](#)

When you configure the ESXi SNMP agent for SNMPv1 and SNMPv2c, the agent supports sending notifications and receiving GET requests.

- [Configure ESXi for SNMP v3](#)

When you configure the ESXi SNMP agent for SNMPv3, the agent supports sending informs and traps. SNMPv3 also provides stronger security than SNMPv1 or SNMPv2c, including key authentication and encryption.

- [Configure the Source of Hardware Events Received by the SNMP Agent](#)

You can configure the ESXi SNMP agent to receive hardware events either from IPMI sensors or CIM indications.

- [Configure the SNMP Agent to Filter Notifications](#)

You can configure the ESXi SNMP agent to filter out notifications if you don't want your SNMP management software to receive those notifications.

- [Configure SNMP Management Client Software](#)

After you have configured a vCenter Server instance or an ESXi host to send traps, configure your management client software to receive and interpret those traps.

## Configure the SNMP Agent for Polling

If you configure the ESXi SNMP agent for polling, it can listen for and respond to requests from SNMP management client systems, such as GET, GETNEXT and GETBULK requests.

By default, the embedded SNMP agent listens on UDP port 161 for polling requests from management systems. You can use the `esxcli system snmp set` command with the `--port` option to configure an alternative port. To avoid conflicting with other services, use a UDP port that is not defined in `/etc/services`.

If you run ESXCLI commands through vCLI, you must supply connection options that specify the target host and login credentials. If you use ESXCLI commands directly on a host using the ESXi Shell, you can use the commands as given without specifying connection options. For more information on connection options see *vSphere Command-Line Interface Concepts and Examples*.

### Prerequisites

Configure the ESXi SNMP agent by using the ESXCLI commands. See *Getting Started with vSphere Command-Line Interfaces* for more information on how to use ESXCLI.

### Procedure

- 1 Run the `esxcli system snmp set` command with the `--port` option to configure the port.

For example, run the following command:

```
esxcli system snmp set --port port
```

Here, *port* is the port the SNMP agent uses to listen for polling requests.

---

**Note** The port you specify must not be already in use by other services. Use IP addresses from the dynamic range, port 49152 and up.

---

- 2 (Optional) If the ESXi SNMP agent is not enabled, run the following command:

```
esxcli system snmp set --enable true
```

## Configure ESXi for SNMPv1 and SNMPv2c

When you configure the ESXi SNMP agent for SNMPv1 and SNMPv2c, the agent supports sending notifications and receiving GET requests.

In SNMPv1 and SNMPv2c, authentication is performed by using community strings. Community strings are namespaces which contain one or more managed objects. This form of authentication does not secure the communication between the SNMP agent and the management system. To secure the SNMP communication in your environment, use SNMPv3.

### Procedure

- 1 [Configure SNMP Communities](#)

To enable the ESXi SNMP agent to send and receive SNMP v1 and v2c messages, you must configure at least one community for the agent.

## 2 Configure the SNMP Agent to Send SNMP v1 or v2c Notifications

You can use the ESXi SNMP agent to send virtual machine and environmental notifications to management systems.

### Configure SNMP Communities

To enable the ESXi SNMP agent to send and receive SNMP v1 and v2c messages, you must configure at least one community for the agent.

An SNMP community defines a group of devices and management systems. Only devices and management systems that are members of the same community can exchange SNMP messages. A device or management system can be a member of multiple communities.

If you run ESXCLI commands through vCLI, you must supply connection options that specify the target host and login credentials. If you use ESXCLI commands directly on a host using the ESXi Shell, you can use the commands as given without specifying connection options. For more information on connection options see *vSphere Command-Line Interface Concepts and Examples*.

#### Prerequisites

Configure the ESXi SNMP agent by using the ESXCLI commands. See *Getting Started with vSphere Command-Line Interfaces* for more information on how to use ESXCLI.

#### Procedure

- ◆ Run the `esxcli system snmp set --communities` option to configure an SNMP community.

For example, to configure public, East, and West network operation centers communities, run the following command:

```
esxcli system snmp set --communities public,eastnoc,westnoc
```

Each time you specify a community with this command, the settings you specify overwrite the previous configuration. To specify multiple communities, separate the community names with a comma.

### Configure the SNMP Agent to Send SNMP v1 or v2c Notifications

You can use the ESXi SNMP agent to send virtual machine and environmental notifications to management systems.

To send SNMP v1/v2c notifications with the SNMP agent, you must configure the target (receiver) unicast address, community, and an optional port. If you do not specify a port, the SNMP agent sends traps to UDP port 162 on the target management system by default.

To configure SNMP v3 traps, see [Configure SNMP v3 Targets](#).

If you run ESXCLI commands through vCLI, you must supply connection options that specify the target host and login credentials. If you use ESXCLI commands directly on a host using the ESXi Shell, you can use the commands as given without specifying connection options. For more information on connection options see *vSphere Command-Line Interface Concepts and Examples*.

### Prerequisites

Configure the ESXi SNMP agent by using the ESXCLI commands. See *Getting Started with vSphere Command-Line Interfaces* for more information on how to use ESXCLI.

### Procedure

- 1 Run the `esxcli system snmp set` command with the `--targets` option:

```
esxcli system snmp set --targets target_address@port/community
```

Here, *target\_address* is the address of the target system, *port* is the port number to send the notifications to, and *community* is the community name.

Each time you specify a target with this command, the settings you specify overwrite all previously specified settings. To specify multiple targets, separate them with a comma.

For example, run the following command for configuring the targets `192.0.2.1@163/westnoc` and `2001:db8::1@163/eastnoc`:

```
esxcli system snmp set --targets 192.0.2.1@163/westnoc,2001:db8::1@163/eastnoc
```

- 2 (Optional) If the ESXi SNMP agent is not enabled, run the following command:

```
esxcli system snmp set --enable true
```

- 3 (Optional) Send a test trap to verify that the agent is configured correctly by running the `esxcli system snmp test` command.

The agent sends a `warmStart` trap to the configured target.

## Configure ESXi for SNMP v3

When you configure the ESXi SNMP agent for SNMPv3, the agent supports sending informs and traps. SNMPv3 also provides stronger security than SNMPv1 or SNMPv2c, including key authentication and encryption.

Inform is a notification that the sender resends up to three times or until the receiver acknowledges the notification.

### Procedure

- 1 [Configure the SNMP Engine ID](#)

Every SNMP v3 agent has an engine ID which serves as a unique identifier for the agent. The engine ID is used with a hashing function to generate keys for authentication and encryption of SNMP v3 messages.



## 2 Configure SNMP Authentication and Privacy Protocols

SNMPv3 optionally supports authentication and privacy protocols.

## 3 Configure SNMP Users

You can configure up to 5 users who can access SNMP v3 information. User names must be no more than 32 characters long.

## 4 Configure SNMP v3 Targets

Configure SNMP v3 targets to allow the ESXi SNMP agent to send SNMP v3 traps and informs.

## Configure the SNMP Engine ID

Every SNMP v3 agent has an engine ID which serves as a unique identifier for the agent. The engine ID is used with a hashing function to generate keys for authentication and encryption of SNMP v3 messages.

If you do not specify an engine ID, when you enable the SNMP agent, an engine ID is automatically generated.

If you run ESXCLI commands through vCLI, you must supply connection options that specify the target host and login credentials. If you use ESXCLI commands directly on a host using the ESXi Shell, you can use the commands as given without specifying connection options. For more information on connection options see *vSphere Command-Line Interface Concepts and Examples*.

### Prerequisites

Configure the ESXi SNMP agent by using the ESXCLI commands. See *Getting Started with vSphere Command-Line Interfaces* for more information on how to use ESXCLI.

### Procedure

- ◆ Run the `esxcli system snmp set` command with the `--engineid` option to configure the SNMP engine ID.

For example, run the following command:

```
esxcli system snmp set --engineid id
```

Here, *id* is the engine ID and it must be a hexadecimal string between 5 and 32 characters long.

## Configure SNMP Authentication and Privacy Protocols

SNMPv3 optionally supports authentication and privacy protocols.

Authentication is used to ensure the identity of users. Privacy allows for encryption of SNMP v3 messages to ensure confidentiality of data. These protocols provide a higher level of security than is available in SNMPv1 and SNMPv2c, which use community strings for security.

Both authentication and privacy are optional. However, you must enable authentication to enable privacy.

The SNMPv3 authentication and privacy protocols are licensed vSphere features and might not be available in some vSphere editions.

If you run ESXCLI commands through vCLI, you must supply connection options that specify the target host and login credentials. If you use ESXCLI commands directly on a host using the ESXi Shell, you can use the commands as given without specifying connection options. For more information on connection options see *vSphere Command-Line Interface Concepts and Examples*.

### Prerequisites

Configure the ESXi SNMP agent by using the ESXCLI commands. See *Getting Started with vSphere Command-Line Interfaces* for more information on how to use ESXCLI.

### Procedure

- 1 (Optional) Run the `esxcli system snmp set` command with the `--authentication` option to configure authentication.

For example, run the following command:

```
esxcli system snmp set --authentication protocol
```

Here, *protocol* must be either **none** (for no authentication), **SHA1**, or **MD5**.

- 2 (Optional) Run the `esxcli system snmp set` command with the `--privacy` option to configure privacy.

For example, run the following command:

```
esxcli system snmp set --privacy protocol
```

Here, *protocol* must be either **none** (for no privacy) or **AES128**.

## Configure SNMP Users

You can configure up to 5 users who can access SNMP v3 information. User names must be no more than 32 characters long.

While configuring a user, you generate authentication and privacy hash values based on the user's authentication and privacy passwords and the SNMP agent's engine ID. If you change the engine ID, the authentication protocol, or the privacy protocol after configuring users, the users are no longer valid and must be reconfigured.

If you run ESXCLI commands through vCLI, you must supply connection options that specify the target host and login credentials. If you use ESXCLI commands directly on a host using the ESXi Shell, you can use the commands as given without specifying connection options. For more information on connection options see *vSphere Command-Line Interface Concepts and Examples*.

### Prerequisites

- Verify that you have configured the authentication and privacy protocols before configuring users.
- Verify that you know the authentication and privacy passwords for each user you plan to configure. Passwords must be at least 7 characters long. Store these passwords in files on the host system.

- Configure the ESXi SNMP agent by using the ESXCLI commands. See *Getting Started with vSphere Command-Line Interfaces* for more information on how to use ESXCLI.

### Procedure

- 1 If you are using authentication or privacy, get the authentication and privacy hash values for the user by running the `esxcli system snmp hash` command with the `--auth-hash` and `--priv-hash` flags.

For example, run the following command:

```
esxcli system snmp hash --auth-hash secret1 --priv-hash secret2
```

Here, *secret1* is the path to the file containing the user's authentication password and *secret2* is the path to the file containing the user's privacy password.

Alternatively, you can pass the `--raw-secret` flag and specify the passwords directly on the command line.

For example, you can run the following command:

```
esxcli system snmp hash --auth-hash authsecret --priv-hash privsecret --raw-secret
```

The produced output might be the following:

```
Authhash: 08248c6eb8b333e75a29ca0af06b224faa7d22d6
Privhash: 232ba5cbe8c55b8f979455d3c9ca8b48812adb97
```

The authentication and privacy hash values are displayed.

- 2 Configure the user by running the `esxcli system snmp set` command with the `--users` flag.

For example, you can run the following command:

```
esxcli system snmp set --users userid/authhash/privhash/security
```

The command accepts the following parameters:

---

Parameter	Description
<i>userid</i>	The user name.
<i>authhash</i>	The authentication hash value.

Parameter	Description
<i>privhash</i>	The privacy hash value.
<i>security</i>	The level of security enabled for that user, which can be <i>auth</i> (for authentication only), <i>priv</i> (for authentication and privacy), or <i>none</i> (for no authentication or privacy).

For example, run the following command to configure user1 for access with authentication and privacy:

```
esxcli system snmp set --users user1/08248c6eb8b333e75a29ca0af06b224faa7d22d6/232ba5cbe8c55b8f979455d3c9ca8b48812adb97/priv
```

You must run the following command to configure user2 for access with no authentication or privacy:

```
esxcli system snmp set --users user2/--/none
```

3 (Optional) Test the user configuration by running the following command:

```
esxcli system snmp test --users username --auth-hash secret1 --priv-hash secret2
```

If the configuration is correct, this command returns the following message: "User *username* validated correctly using engine id and security level: *protocols*". Here, *protocols* indicates the security protocols configured.

## Configure SNMP v3 Targets

Configure SNMP v3 targets to allow the ESXi SNMP agent to send SNMP v3 traps and informs.

SNMP v3 allows for sending both traps and informs. An inform message is a type of a message that the sender resends a maximum of three times. The sender waits for 5 seconds between each attempt, unless the message is acknowledged by the receiver.

You can configure a maximum of three SNMP v3 targets, in addition to a maximum of three SNMP v1/v2c targets.

To configure a target, you must specify a hostname or IP address of the system that receives the traps or informs, a user name, a security level, and whether to send traps or informs. The security level can be either **none** (for no security), **auth** (for authentication only), or **priv** (for authentication and privacy).

If you run ESXCLI commands through vCLI, you must supply connection options that specify the target host and login credentials. If you use ESXCLI commands directly on a host using the ESXi Shell, you can use the commands as given without specifying connection options. For more information on connection options see *vSphere Command-Line Interface Concepts and Examples*.

### Prerequisites

- Ensure that the users who access the traps or informs are configured as SNMP users for both the ESXi SNMP agent and the target management system.
- If you are configuring informs, you need the engine ID for the SNMP agent on the remote system that receives the inform message.

- Configure the ESXi SNMP agent by using the ESXCLI commands. See *Getting Started with vSphere Command-Line Interfaces* for more information on how to use ESXCLI.

## Procedure

- 1 (Optional) If you are configuring informs, configure the remote users by running the `esxcli system snmp set` command with the `--remote-users` option.

For example, run the following command:

```
esxcli system snmp set --remote-users userid/auth-protocol/auth-hash/priv-protocol/priv-hash/engine-id
```

The command accepts the following parameters:

Parameter	Description
<i>userid</i>	The user name.
<i>auth-protocol</i>	The authentication protocol, <b>none</b> (for no authentication), <b>MD5</b> , or <b>SHA1</b> .
<i>auth-hash</i>	The authentication hash or – if the authentication protocol is <b>none</b> .
<i>priv-protocol</i>	The privacy protocol, <b>AES128</b> , or <b>none</b> .
<i>priv-hash</i>	The privacy hash, or – if the privacy protocol is <b>none</b> .
<i>engine-id</i>	The engine ID of the SNMP agent on the remote system that receives the inform message.

- 2 Run the `esxcli system snmp set` command with the `--v3targets` option.

For example, run the following command:

```
esxcli system snmp set --v3targets hostname@port/userid/secLevel/message-type
```

The parameters of the command are as follows.

Parameter	Description
<i>hostname</i>	The host name or IP address of the management system that receives the traps or informs.
<i>port</i>	The port on the management system that receives the traps or informs. If you do not specify a port, the default port, 162, is used.
<i>userid</i>	The user name.
<i>secLevel</i>	The level of authentication and privacy you have configured. Use <b>auth</b> if you have configured authentication only, <b>priv</b> if you have configured both authentication and privacy, and <b>none</b> if you have configured neither.
<i>message-type</i>	The type of the messages received by the management system. Use <b>trap</b> or <b>inform</b> .

- 3 (Optional) If the ESXi SNMP agent is not enabled, run the following command:

```
esxcli system snmp set --enable true
```

- 4 (Optional) Send a test notification to verify that the agent is configured correctly by running the `esxcli system snmp test` command.

The agent sends a `warmStart` notification to the configured target.

## Configure the Source of Hardware Events Received by the SNMP Agent

You can configure the ESXi SNMP agent to receive hardware events either from IPMI sensors or CIM indications.

IPMI sensors are used for hardware monitoring in ESX/ESXi 4.x and earlier. The conversion of CIM indications to SNMP notifications is available in ESXi 5.0 and later.

If you run ESXCLI commands through vCLI, you must supply connection options that specify the target host and login credentials. If you use ESXCLI commands directly on a host using the ESXi Shell, you can use the commands as given without specifying connection options. For more information on connection options see *vSphere Command-Line Interface Concepts and Examples*.

### Prerequisites

Configure the ESXi SNMP agent by using the ESXCLI commands. See *Getting Started with vSphere Command-Line Interfaces* for more information on how to use ESXCLI.

### Procedure

- 1 Run the `esxcli system snmp set --hwsrc source` command to configure the source for hardware events.

Here, *source* is **sensors** or **indications**, for hardware event received from IPMI sensors or CIM indications respectively.

- 2 (Optional) If the ESXi SNMP agent is not enabled, run the following command:

```
esxcli system snmp set --enable true
```

## Configure the SNMP Agent to Filter Notifications

You can configure the ESXi SNMP agent to filter out notifications if you don't want your SNMP management software to receive those notifications.

### Prerequisites

Configure the ESXi SNMP agent by using the ESXCLI commands. See *Getting Started with vSphere Command-Line Interfaces* for more information on how to use ESXCLI.

### Procedure

- 1 Run the `esxcli system snmp set` command to filter notifications:

```
esxcli system snmp set --notraps oid_list
```

Here, *oid\_list* is a list of OIDs for the notifications to filter, separated by commas. This list replaces any OIDs that were previously specified using this command.

For example, to filter out coldStart (OID 1.3.6.1.4.1.6876.4.1.1.0) and warmStart (OID 1.3.6.1.4.1.6876.4.1.1.1) traps, run the following command:

```
esxcli system snmp set --notraps 1.3.6.1.4.1.6876.4.1.1.0,1.3.6.1.4.1.6876.4.1.1.1
```

- 2 (Optional) If the ESXi SNMP agent is not enabled, run the following command:

```
esxcli system snmp set --enable true
```

The traps identified by the specified OIDs are filtered out of the output of the SNMP agent, and are not sent to SNMP management software.

### What to do next

To clear all notification filters, run the `esxcli system snmp set --notraps reset` command.

## Configure SNMP Management Client Software

After you have configured a vCenter Server instance or an ESXi host to send traps, configure your management client software to receive and interpret those traps.

To configure your management client software, specify the communities for the managed device, configure the port settings, and load the VMware MIB files. See the documentation for your management system for specific instructions for these steps.

### Prerequisites

To complete this task, download the VMware MIB files from the VMware website:

<http://communities.vmware.com/community/developer/managementapi>. On the Web page, search Downloading MIB modules.

### Procedure

- 1 In your management software, specify the vCenter Server instance or ESXi host as an SNMP-based managed device.
- 2 If you are using SNMPv1 or SNMPv2c, set up appropriate community names in the management software.  
  
These names must correspond to the communities set for the SNMP agent on the vCenter Server instance or ESXi host.
- 3 If you are using SNMPv3, configure users and authentication and privacy protocols to match those configured on the ESXi host.
- 4 If you configured the SNMP agent to send traps to a port on the management system other than the default UDP port 162, configure the management client software to listen on the port you configured.

- 5 Load the VMware MIBs into the management software so you can view the symbolic names for vCenter Server or the host variables.

To prevent lookup errors, load these MIB files in the following order before loading other MIB files:

- a VMWARE-ROOT-MIB.mib
- b VMWARE-TC-MIB.mib
- c VMWARE-PRODUCTS-MIB.mib

The management software can now receive and interpret traps from vCenter Server or ESXi hosts.

## SNMP Diagnostics

You can use SNMP tools to diagnose configuration problems.

- Run the `esxcli system snmp test` command from the vSphere CLI set to prompt the SNMP agent to send a test `warmStart` trap.
- Run the `esxcli system snmp get` command to display the current configuration of the SNMP agent.
- The `SNMPv2-MIB.mib` file provides several counters to aid in debugging SNMP problems. See [SNMPv2 Diagnostic Counters](#).
- The `VMWARE-AGENTCAP-MIB.mib` file defines the capabilities of the VMware SNMP agents by product version. Use this file to determine if the SNMP functionality that you want to use is supported.

## Monitor Guest Operating Systems with SNMP

You can use SNMP to monitor guest operating systems or applications running in virtual machines.

The virtual machine uses its own virtual hardware devices. Do not install agents in the virtual machines that are intended to monitor physical hardware.

### Procedure

- ◆ Install the SNMP agents you normally would use for that purpose in the guest operating systems.

## VMware MIB Files

VMware MIB files define the information provided by ESXi hosts and vCenter Server to SNMP management software.

You can download these MIB files from

<http://communities.vmware.com/community/developer/forums/managementapi#SNMP-MIB>. On the Web page, search for Downloading MIB modules.

The table VMware MIB Files lists the MIB files provided by VMware and describes the information that each file provides.



**Table 10-1. VMware MIB Files**

<b>MIB File</b>	<b>Description</b>
VMWARE-ROOT-MIB.mib	Contains VMware's enterprise OID and top-level OID assignments.
VMWARE-AGENTCAP-MIB.mib	Defines the capabilities of the VMware agents by product versions. This file is optional and might not be supported by all management systems.
VMWARE-CIMOM-MIB.mib	Defines variables and trap types used to report on the state of the CIM Object Management subsystem.
VMWARE-ENV-MIB.mib	Defines variables and trap types used to report on the state of physical hardware components of the host computer. Enables conversion of CIM indications to SNMP traps.
VMWARE-OBSOLETE-MIB.mib	For use with versions of ESX/ESXi prior to 4.0. Defines OIDs that have been made obsolete to maintain backward compatibility with earlier versions of ESX/ESXi. Includes variables formerly defined in the files VMWARE-TRAPS-MIB.mib and VMWARE-VMKERNEL-MIB.mib.
VMWARE-PRODUCTS-MIB.mib	Defines OIDs to uniquely identify each SNMP agent on each VMware platform by name, version, and build platform.
VMWARE-RESOURCES-MIB.mib	Defines variables used to report information on resource usage of the VMkernel, including physical memory, CPU, and disk utilization.
VMWARE-SYSTEM-MIB.mib	The VMWARE-SYSTEM-MIB.mib file is obsolete. Use the SNMPv2-MIB to obtain information from sysDescr.0 and sysObjec ID.0.
VMWARE-TC-MIB.mib	Defines common textual conventions used by VMware MIB files.
VMWARE-VC-EVENTS-MIB.mib	Defines traps sent by vCenter Server. Load this file if you use vCenter Server to send traps.
VMWARE-VMINFO-MIB.mib	Defines variables for reporting information about virtual machines, including virtual machine traps.

The table Other MIB Files lists MIB files included in the VMware MIB files package that are not created by VMware. These can be used with the VMware MIB files to provide additional information.

**Table 10-2. Other MIB Files**

<b>MIB File</b>	<b>Description</b>
ENTITY-MIB.mib	Allows the description of relationships between physical entities and logical entities managed by the same SNMP agent. See RFC 4133 for more information.
HOST-RESOURCES-MIB.mib	Defines objects that are useful for managing host computers.
HOST-RESOURCES-TYPES.mib	Defines storage, device, and filesystem types for use with HOST-RESOURCES-MIB.mib.
IEEE8021-BRIDGE-MIB	Defines objects for managing devices that support IEEE 802.1D.
IEEE8023-LAG-MIB	Defines objects for managing devices that support IEEE 802.3ad link aggregation.
IEEE8021-Q-BRIDGE-MIB	Defines objects for managing Virtual Bridged Local Area Networks.
IF-MIB.mib	Defines attributes related to physical NICs on the host system.

**Table 10-2. Other MIB Files (Continued)**

MIB File	Description
IP-MIB.mib	Defines objects for managing implementations of the Internet Protocol (IP) in an IP version-independent manner.
IP-FORWARD-MIB.mib	Defines objects for managing IP forwarding.
LLDP-V2-MIB.mib	Defines objects for managing devices using Linked Layer Discovery Protocol (LLDP).
SNMPv2-CONF.mib	Defines conformance groups for MIBs.
SNMPv2-MIB.mib	Defines the SNMP version 2 MIB objects.
SNMPv2-SMI.mib	Defines the Structure of Management Information for SNMP version 2.
SNMPv2-TC.mib	Defines textual conventions for SNMP version 2.
TCP-MIB.mib	Defines objects for managing devices using the TCP protocol.
UDP-MIB.mib	Defines objects for managing devices using the UDP protocol.

## SNMPv2 Diagnostic Counters

The `SNMPv2-MIB.mib` file provides a number of counters to aid in debugging SNMP problems.

[Table 10-3](#) lists some of these diagnostic counters.

**Table 10-3. Diagnostic Counters from SNMPv2-MIB**

Variable	ID Mapping	Description
<code>snmpInPkts</code>	<code>snmp 1</code>	The total number of messages delivered to the SNMP entity from the transport service.
<code>snmpInBadVersions</code>	<code>snmp 3</code>	The total number of SNMP messages that were delivered to the SNMP entity and were for an unsupported SNMP version.
<code>snmpInBadCommunityNames</code>	<code>snmp 4</code>	The total number of community-based SNMP messages delivered to the SNMP entity that used an invalid SNMP community name.
<code>snmpInBadCommunityUses</code>	<code>snmp 5</code>	The total number of community-based SNMP messages delivered to the SNMP entity that represented an SNMP operation that was not allowed for the community named in the message.
<code>snmpInASNParseErrs</code>	<code>snmp 6</code>	The total number of ASN.1 or BER errors encountered by the SNMP entity when decoding received SNMP messages.
<code>snmpEnableAuthenTraps</code>	<code>snmp 30</code>	Indicates whether the SNMP entity is permitted to generate authenticationFailure traps. The value of this object overrides any configuration information. It therefore provides a means of disabling all authenticationFailure traps.

**Table 10-3. Diagnostic Counters from SNMPv2-MIB (Continued)**

<b>Variable</b>	<b>ID Mapping</b>	<b>Description</b>
snmpSilentDrops	snmp 31	The total number of Confirmed Class PDUs delivered to the SNMP entity that were silently dropped because the size of a reply containing an alternate Response Class PDU with an empty variable-bindings field was greater than either a local constraint or the maximum message size associated with the originator of the request.
snmpProxyDrops	snmp 32	The total number of Confirmed Class PDUs delivered to the SNMP entity that were silently dropped because the transmission of the message to a proxy target failed in a manner other than a time-out such that no Response Class PDU could be returned.

# System Log Files

In addition to lists of events and alarms, vSphere components generate assorted logs. These logs contain additional information about activities in your vSphere environment.

This chapter includes the following topics:

- [View System Log Entries](#)
- [View System Logs on an ESXi Host](#)
- [System Logs](#)
- [Export System Log Files](#)
- [ESXi Log Files](#)
- [Upload Logs Package to a VMware Service Request](#)
- [Configure Syslog on ESXi Hosts](#)
- [Configuring Logging Levels for the Guest Operating System](#)
- [Collecting Log Files](#)
- [Viewing Log Files with the Log Browser](#)

## View System Log Entries

You can view the system logs generated by vSphere components. These instructions apply only to vCenter Server management nodes.

### Procedure

- 1 In the vSphere Web Client, navigate to a vCenter Server.
- 2 From the Monitor tab, click **System Logs**.
- 3 From the drop-down menu, select the log and entry you want to view.

## View System Logs on an ESXi Host

You can use the direct console interface to view the system logs on an ESXi host. These logs provide information about system operational events.

**Procedure**

1 From the direct console, select **View System Logs**.

2 Press a corresponding number key to view a log.

vCenter Server agent (vpxa) logs appear if the host is managed by vCenter Server.

3 Press Enter or the spacebar to scroll through the messages.

4 (Optional) Perform a regular expression search.

a Press the slash key (/).

b Type the text to find.

c Press Enter

The found text is highlighted on the screen.

5 Press q to return to the direct console.

## System Logs

VMware technical support might request several files to help resolve any issues you might have with the product. This section describes the types and locations of log files found on ESXi hosts and vSphere Web Client.

### ESXi System Logs

You might need the ESXi system log files to resolve technical issues.

The ESXi system logs can be found in the `/var/run/log` directory.

### vSphere Web Client Logs

You might need the vSphere Web Client system log files to resolve technical issues.

Depending on, whether you use a vCenter Server instance that runs on Windows, or a vCenter Server Appliance, the vSphere Web Client system logs can be found in the location listed in the table.

**Table 11-1. Location of vSphere Web Client Logs**

<b>vCenter Server System</b>	<b>Location</b>
vCenter Server that runs on Windows	<code>C:\ProgramData\VMware\vCenterServer\logs\vsphere-client\logs</code>
vCenter Server Appliance	<code>/var/log/vmware/vsphere-client/logs</code>

The main vSphere Web Client log file is `vsphere_client_virgo.log`.

## Export System Log Files

When the vSphere Web Client is connected to vCenter Server, you can select hosts from which to download system log files.

To save diagnostic data for ESXi hosts and vCenter Server, the vSphere Web Client must be connected to the vCenter Server system.

Required privileges:

- To view diagnostic data: **Read-Only User**
- To manage diagnostic data: **Global.Diagnostics**

### Procedure

- 1 In the inventory, navigate to a vCenter Server instance, and click the **Monitor** tab.
- 2 Click **System Logs**.
- 3 Click **Export System Logs**.
- 4 If you are connected to vCenter Server, select the object for which you want to export data.  
Selecting an object selects all of its child objects.
- 5 If you are connected to vCenter Server, select **Include information from vCenter Server and vSphere Web Client** to download vCenter Server and vSphere Web Client log files and host log files, and click **Next**.
- 6 If the selected host supports manifest driven exports of system log files, select the system log files to collect. Select the specific system log files to download.  
If the host does not support manifest exports of log files, all system log files are exported.
- 7 Select **Gather performance data** to include performance data information in the log files.  
You can update the duration and interval time you want the data collected.
- 8 (Optional) Select to apply a password for encrypted core dumps to the support package.  
You can make that password available to your support representative in a secure channel.  
If only some of the host in your environment use encryption, some of the files in the package are encrypted.
- 9 Click **Finish**.
- 10 Specify the location to which to save the log files.  
The host or vCenter Server generates a .zip file containing the log files.
- 11 Click **Save**.  
The **Recent Tasks** panel shows the Generate diagnostic bundles task in progress.

The Downloading Log Bundles dialog box appears when the Generating Diagnostic Bundle task is finished. The download status of each bundle appears in the dialog box.

Some network errors can cause download failures. When you select an individual download in the dialog box, the error message for that operation appears under the name and location of the log bundle file.

12 Verify the information in the Summary and click **Finish** to download the log files.

Diagnostic bundles containing log files for the specified objects are downloaded to the location specified.

#### What to do next

[Upload Logs Package to a VMware Service Request.](#)

## ESXi Log Files

Log files are an important component of troubleshooting attacks and obtaining information about breaches. Logging to a secure, centralized log server can help prevent log tampering. Remote logging also provides a long-term audit record.

To increase the security of the host, take the following measures

- Configure persistent logging to a datastore. By default, the logs on ESXi hosts are stored in the in-memory file system. Therefore, they are lost when you reboot the host, and only 24 hours of log data is stored. When you enable persistent logging, you have a dedicated activity record for the host.
- Remote logging to a central host allows you to gather log files on a central host. From that host, you can monitor all hosts with a single tool, do aggregate analysis, and search log data. This approach facilitates monitoring and reveals information about coordinated attacks on multiple hosts.
- Configure the remote secure syslog on ESXi hosts by using a CLI such as vCLI or PowerCLI, or by using an API client.
- Query the syslog configuration to make sure that the syslog server and port are valid.

See the *vSphere Monitoring and Performance* documentation for information about syslog setup, and for additional information on ESXi log files.

## Upload Logs Package to a VMware Service Request

If you already have a Service Request ID with VMware, you can use the vSphere Web Client to upload the system log bundles directly to your service request.

#### Prerequisites

Request a Service Request ID from VMware Technical Support.

#### Procedure

1 In the vSphere Web Client, navigate to **Administration**.

- 2 Click **Support**, and click **Upload File to Service Request**.

An Upload File to Service Request dialog box opens.

- 3 Enter your Service Request ID with VMware.
- 4 Click **Choose File**, and select the log bundle you want to attach to your service request with VMware, and click **OK**.
- 5 If you protected your support package with a password, provide the password to VMware Technical Support by using a secure channel.

The log bundle is sent to your service request.

## Configure Syslog on ESXi Hosts

You can use the vSphere Web Client or the `esxcli system syslog vCLI` command to configure the syslog service.

For information about using the `esxcli system syslog` command and other vCLI commands, see *Getting Started with vSphere Command-Line Interfaces*.

### Procedure

- 1 In the vSphere Web Client inventory, select the host.
- 2 Click **Configure**.
- 3 Under System, click **Advanced System Settings**.
- 4 Filter for **syslog**.
- 5 To set up logging globally, select the setting to change and click **Edit**.

Option	Description
<b>Syslog.global.defaultRotate</b>	Maximum number of archives to keep. You can set this number globally and for individual subloggers.
<b>Syslog.global.defaultSize</b>	Default size of the log, in KB, before the system rotates logs. You can set this number globally and for individual subloggers.
<b>Syslog.global.LogDir</b>	Directory where logs are stored. The directory can be located on mounted NFS or VMFS volumes. Only the <code>/scratch</code> directory on the local file system is persistent across reboots. Specify the directory as <code>[datastorename] path_to_file</code> , where the path is relative to the root of the volume backing the datastore. For example, the path <code>[storage1] /systemlogs</code> maps to the path <code>/vmfs/volumes/storage1/systemlogs</code> .



Option	Description
<b>Syslog.global.logDirUnique</b>	Selecting this option creates a subdirectory with the name of the ESXi host under the directory specified by <b>Syslog.global.LogDir</b> . A unique directory is useful if the same NFS directory is used by multiple ESXi hosts.
<b>Syslog.global.LogHost</b>	Remote host to which syslog messages are forwarded and port on which the remote host receives syslog messages. You can include the protocol and the port, for example, <code>ssl://hostName1:1514</code> . UDP (default), TCP, and SSL are supported. The remote host must have syslog installed and correctly configured to receive the forwarded syslog messages. See the documentation for the syslog service installed on the remote host for information on configuration.

6 (Optional) To overwrite the default log size and log rotation for any of the logs.

- a Click the name of the log that you want to customize.
- b Click **Edit** and enter the number of rotations and the log size you want.

7 Click **OK**.

Changes to the syslog options take effect immediately.

## Configuring Logging Levels for the Guest Operating System

Virtual machines can write support and troubleshooting information into a virtual machine log file stored on a VMFS volume. The default settings for virtual machines are appropriate for most situations.

If your environment relies heavily on using vMotion, or if the defaults do not seem suitable for other reasons, you can modify the logging settings for virtual machine guest operating systems.

New log file creation happens as follows:

- Each time you power on or resume a virtual machine, and each time you migrate a virtual machine with vMotion, a new log file is created.
- Each time an entry is written to the log, the size of the log is checked. If `vmx.log.rotateSize` is set to a nondefault value, and the size is over the limit, the next entry is written to a new log. If the maximum number of log files exists, the oldest log file is deleted.

The default for `vmx.log.rotateSize` is zero (0), which means new logs are created during power-on, resume, and so on. You can ensure the new log file creation happens more frequently by limiting the maximum size of the log files with the `vmx.log.rotateSize` configuration parameter.

VMware recommends saving 10-log files, each one limited to no less than 2MB. These values are large enough to capture sufficient information to debug most problems. If you need logs for a longer time span, you can set `vmx.log.keepOld` to 20.

## Change the Number of Virtual Machine Log Files

You can change the number of the log files for all virtual machines on an ESXi host or for individual virtual machines.

This procedure discusses limiting the virtual machine log file number on an individual virtual machine.

To limit the number of log files for *all* virtual machines on a host, edit the `/etc/vmware/config` file. If the `vmx.log.KeepOld` property is not defined in the file, you can add it. For example, to keep ten log files for each virtual machine, add the following to `/etc/vmware/config`:

```
vmx.log.KeepOld = "10"
```

You can use a PowerCLI script to change this property on all the virtual machines on a host.

You can use the `log.KeepOld` parameter to affect all log files, not just the virtual machine log files.

### Prerequisites

Turn off the virtual machine.

### Procedure

- 1 Log in to a vCenter Server system using the vSphere Web Client and find the virtual machine.
  - a In the Navigator, select **VMs and Templates**.
  - b Find the virtual machine in the hierarchy.
- 2 Right-click the virtual machine and click **Edit Settings**.
- 3 Select **VM Options**.
- 4 Click **Advanced** and click **Edit Configuration**.
- 5 Add or edit the `vmx.log.KeepOld` parameter to the number of files to keep for this virtual machine.
 

For example, to keep 20 log files and begin deleting the oldest files as new ones are created, enter **20**.
- 6 Click **OK**.

## Control When to Switch to New Virtual Machine Log Files

The `vmx.log.rotateSize` parameter specifies the log file size at which the switch to new log files happens for the logs for individual virtual machines. Use this parameter together with the `vmx.log.KeepOld` parameter to ensure acceptable log file sizes without losing critical logging information.

The `vmx.log.KeepOld` parameter determines how many virtual machine log file instances the ESXi host retains before overwriting the first log file. The default value of `vmx.log.KeepOld` is 10, a suitable number to properly log complex operations such as vMotion. You must increase this number significantly when you change the value of `vmx.log.rotateSize`.

This procedure discusses changing the virtual machine rotate size on an individual virtual machine.

To limit the rotate size for *all* virtual machines on a host, edit the `/etc/vmware/config` file. If the `vmx.log.KeepOld` property is not defined in the file, you can add it. You can use a PowerCLI script to change this parameter for selected virtual machines on a host.

You can use the `log.rotateSize` parameter to affect all log files, not just the virtual machine log files.

You can change the value of `vmx.log.rotateSize` for all virtual machines from the vSphere Web Client or by using a PowerCLI script.

### Prerequisites

Turn off the virtual machine.

### Procedure

- 1 Log in to a vCenter Server system using the vSphere Web Client and find the virtual machine.
  - a In the Navigator, select **VMs and Templates**.
  - b Find the virtual machine in the hierarchy.
- 2 Right-click the virtual machine and click **Edit Settings**.
- 3 Select **VM Options**.
- 4 Click **Advanced** and click **Edit Configuration**.
- 5 Add or edit the `vmx.log.rotateSize` parameter to the maximum file size before log information is added to a new file.

Alternatively, you can add or edit the first log file if you have more log files than the `vmx.log.keepOld` parameter specifies.

Specify the size in bytes.

- 6 Click **OK**.

## Collecting Log Files

VMware technical support might request several files to help resolve technical issues. The following sections describe script processes for generating and collecting some of these files.

### Set Verbose Logging

You can set the verbose log file specification

You can only set verbose logging for vpxd logs.

### Procedure

- 1 Select **Administration > vCenter Server Settings**.
- 2 Select **Logging Options**.
- 3 Select **Verbose** from the pop-up menu.
- 4 Click **OK**.

## Collect vSphere Log Files

You can collect vSphere log files in to a single location.

### Procedure

- ◆ View the log file using one of the following methods.

Task	Action
View the <code>viclient-*.log</code> file	Change to the directory, <code>%temp%</code> .
Download the log bundle from vSphere Web Client connected to a vCenter Server system	To download the log bundle, do the following: <ol style="list-style-type: none"> <li>Select <b>Administration &gt; System Configuration</b>.</li> <li>From the Objects tab, select <b>Actions &gt; Export Support Bundles...</b></li> </ol> The log bundle is generated as a <code>.zip</code> file. By default, the <code>vpxd</code> logs within the bundle are compressed as <code>.tgz</code> files. You must use <code>gunzip</code> to uncompress these files.
Generate vCenter Server log bundles from a vCenter Server system	Select <b>Start &gt; Programs &gt; VMware &gt; Generate vCenter Server log bundle</b> . You can use this to generate vCenter Server log bundles even when you are unable to connect to the vCenter Server using the vSphere Web Client. The log bundle is generated as a <code>.zip</code> file. By default, the <code>vpxd</code> logs within the bundle are compressed as <code>.tgz</code> files. You must use <code>gunzip</code> to uncompress these files.

## Collect ESXi Log Files

You can collect and package all relevant ESXi system and configuration information, a well as ESXi log files. This information can be used to analyze the problems.

### Procedure

- ◆ Run the following script on the ESXi Shell: `/usr/bin/vm-support`

The resulting file has the following format: `esx-date-unique-xnumber.tgz`

## ESXi Log File Locations

ESXi records host activity in log files, using a syslog facility.

Component	Location	Purpose
VMkernel	<code>/var/log/vmkernel.log</code>	Records activities related to virtual machines and ESXi.
VMkernel warnings	<code>/var/log/vmkwarning.log</code>	Records activities related to virtual machines.
VMkernel summary	<code>/var/log/vmksummary.log</code>	Used to determine uptime and availability statistics for ESXi (comma separated).

Component	Location	Purpose
ESXi host agent log	<code>/var/log/hostd.log</code>	Contains information about the agent that manages and configures the ESXi host and its virtual machines.
vCenter agent log	<code>/var/log/vpxa.log</code>	Contains information about the agent that communicates with vCenter Server (if the host is managed by vCenter Server).
Shell log	<code>/var/log/shell.log</code>	Contains a record of all commands typed into the ESXi Shell as well as shell events (for example, when the shell was enabled).
Authentication	<code>/var/log/auth.log</code>	Contains all events related to authentication for the local system.
System messages	<code>/var/log/syslog.log</code>	Contains all general log messages and can be used for troubleshooting. This information was formerly located in the messages log file.
Virtual machines	The same directory as the affected virtual machine's configuration files, named <code>vmware.log</code> and <code>vmware*.log</code> . For example, <code>/vmfs/volumes/datastore/virtual machine/vmware.log</code>	Contains virtual machine power events, system failure information, tools status and activity, time sync, virtual hardware changes, vMotion migrations, machine clones, and so on.

## Configure Log Filtering on ESXi Hosts

The log filtering capability lets you modify the logging policy of the syslog service that is running on an ESXi host. You can create log filters to reduce the number of repetitive entries in the ESXi logs and to blacklist specific log events entirely.

Log filters affect all log events that are processed by the ESXi host `vmsyslogd` daemon, whether they are recorded to a log directory or to a remote syslog server.

When you create a log filter, you set a maximum number of log entries for the log messages. The logs messages are generated by one or more specified system components and that match a specified phrase. You must enable the log filtering capability and reload the syslog daemon to activate the log filters on the ESXi host.

---

**Important** Setting a limit to the amount of logging information, restricts your ability to troubleshoot potential system failures properly. If a log rotate occurs after the maximum number of log entries is reached, you might lose all instances of a filtered message.

---

### Procedure

- 1 Log in to the ESXi Shell as root.

- In the `/etc/vmware/logfilters` file, add the following entry to create a log filter.

```
numLogs | ident | logRegexp
```

where:

- `numLogs` sets the maximum number of log entries for the specified log messages. After reaching this number, the specified log messages are filtered and ignored. Use `0` to filter and ignore all the specified log messages.
- `ident` specifies one or more system components to apply the filter to the log messages that these components generate. For information about the system components that generate log messages, see the values of the `idents` parameters in the `syslog` configuration files. The files are located in the `/etc/vmsyslog.conf.d` directory. Use a comma-separated list to apply a filter to more than one system component. Use `*` to apply a filter to all system components.
- `logRegexp` specifies a case-sensitive phrase with Python regular expression syntax to filter the log messages by their content.

For example, to set a limit of maximum two-log entries from the `hostd` component for messages that resemble the `SOCKET connect failed, error 2: No such file or directory` phrase with any error number, add the following entry:

```
2 | hostd | SOCKET connect failed, error .*: No such file or directory
```

---

**Note** A line starting with `#` denotes a comment and the rest of the line is ignored.

---

- In the `/etc/vmsyslog.conf` file, add the following entry to enable the log filtering capability.

```
enable_logfilters = true
```

- Run the `esxcli system syslog reload` command to reload the `syslog` daemon and apply the configuration changes.

## Turn Off Compression for vpxd Log Files

By default, vCenter Server `vpxd` log files are rolled up and compressed into `.gz` files. You can turn off this setting to leave the `vpxd` logs uncompressed.

### Procedure

- Log in to the vCenter Server using the vSphere Web Client.
- Select **Administration > vCenter Server Settings**.
- Select **Advanced Settings**.
- In the **Key** text box, type `log.compressOnRoll`.
- In the **Value** text box, type `false`.
- Click **Add**, and click **OK**.

## ESXi VMkernel Files

If the VMkernel fails, an error message appears and then the virtual machine reboots. If you specified a VMware core dump partition when you configured your virtual machine, the VMkernel also generates a core dump and error log.

More serious problems in the VMkernel can freeze the machine without an error message or core dump.

## Viewing Log Files with the Log Browser

Log browser is a plug-in to the vSphere Web Client, and is part of the vCenter Server installation package. To use the log browser, you must deploy the log browser plug-in.

If you deploy the log browser plug-in, you can view, search, and export one or more vCenter Server and ESXi log files at a time using the log browser. You can also export, manage, and view different log types.

---

**Note** The log browser cannot be used for the Platform Services Controller in the vCenter Server Appliance or vCenter Server on Windows.

---

## Enable the Log Browser Plug-In on the vCenter Server Appliance

In vSphere 6.5 the log browser plug-in is part of the vCenter Server Appliance, but is not enabled by default and therefore the vSphere Web Client does not display it. You can manually deploy the log browser plug-in on your vCenter Server Appliance.

### Prerequisites

- Verify you have administrative rights to access the vCenter Server Appliance.

### Procedure

- 1 Log in to the vCenter Server Appliance Bash Shell as an Administrator.
- 2 Navigate to the location of the log browser manifest file.

The default location is: `/usr/lib/vmware-vsphere-client/plugin-packages/logbrowser`

- 3 Rename the file `plugin-package.xml.unused` to `plugin-package.xml`, and save it.
- 4 From the vSphere Web Client, restart the VMware Service Lifecycle Manager API service.

For more information how to restart services in the vSphere Web Client, see *vCenter Server and Host Management* documentation.

The **Log Browser** tab appears under the **Monitor** tab in the vSphere Web Client.

## Enable the Log Browser Plug-In on a vCenter Server Instance That Runs on Windows

In vSphere 6.5 the log browser plug-in is part of the vCenter Server installation package, but is not enabled by default and therefore the vSphere Web Client does not display it. You can manually deploy the log browser plug-in on your vCenter Server system that runs on Windows.

### Prerequisites

- Verify you have administrative privileges to access the Windows machine where vCenter Server runs.

### Procedure

- 1 Log in as an administrator to the Windows machine where vCenter Server runs.
- 2 Navigate to the location of the log browser manifest file.

The location of this file in Windows Server 2008/2012 is

```
C:\ProgramData\VMware\vCenterServer\runtime\vsphere-client\plugin-packages\logbrowser.
```

- 3 Rename the file `plugin-package.xml.unused` to `plugin-package.xml`, and save it.
- 4 From the vSphere Web Client, restart the VMware Service Lifecycle Manager API service.

For more information how to restart services in the vSphere Web Client, see *vCenter Server and Host Management* documentation.

The vCenter Server instance restarts.

The **Log Browser** tab appears under the **Monitor** tab in the vSphere Web Client.

## Retrieve Logs

When you retrieve logs for a host or vCenter Server, you can use these logs to view, search, filter, and compare with other system logs.

### Procedure

- 1 Navigate to the host or vCenter Server that contain the logs you want to retrieve.
- 2 Click the **Monitor** tab.
- 3 Click **Log Browser**.
- 4 (Optional) If no logs for the host or vCenter Server are available, click **Retrieve now** to retrieve the logs for that object.

The retrieved logs are based on a current snapshot of the system. Retrieving logs can take a few minutes. You can perform other tasks while the logs are being retrieved.

- 5 (Optional) Click **Refresh** to retrieve newer logs.
- 6 Select the type of log you want to browse.



The log displays in the browser.

## Search Log Files

You can search the log files by text or by time.

### Prerequisites

If the log is unavailable, you must retrieve it. See [Retrieve Logs](#).

### Procedure

- 1 Navigate to the Log Browser.
- 2 Select the type of log you want to browse.
- 3 In the **Actions** menu, select either **Find by Text** or **Find by Time**.
- 4 In the search area at the bottom of the Log Browser type the text or select the time you want to search.
- 5 Click **Next** to view the next line containing the text or time searched or **Previous** to view the previous line.

The Log Browser displays the line (in the 3rd row) that contains the text or time you searched.

## Filter Log Files

The Log Browser displays filtered searches.

### Procedure

- 1 Navigate to the Log Browser and select a log file to browse.
- 2 Select the number of **Adjacent** lines you want the Log Browser to display.
- 3 Type the text that you want filtered in the search box.

When displaying adjacent lines, groups of consecutive entries are highlighted with a different background color.

The log browser displays the lines in the log that contain the word you typed, with the number of adjacent (before and after) lines.

## Create Advanced Log Filters

You can create and save log file filters to your local system.

### Procedure

- 1 Navigate to the Log Browser and retrieve a log file object.
- 2 Click **Advanced filter**.
- 3 Enter the conditions you want included in the filter.

- 4 Type a filter name.
- 5 Click **Save** to save the filter.

The filter is saved on the vSphere Web Client server and is available the next time you start the vSphere Web Client.

- 6 Click **Filter** to view the results in the Log Browser.

#### What to do next

You can load saved filters from your local system by clicking **Save to local system**. The filters are saved in XML format. You can also load filters from an XML file from your local system by clicking **Load from local system**.

## Adjust Log Times

You might want to adjust the times in the log files to a different time zone or to compare multiple log files.

#### Procedure

- 1 Navigate to the Log Browser and retrieve a log file from an object.
- 2 In the **Actions** menu, select **Adjust by Time**.
- 3 Select **Add** or **Subtract** and adjust the **days**, **hours**, **minutes**, **seconds**, or **milliseconds** from the original time stamps in the log.

The adjusted time stamp appears.

- 4 (Optional) Click **Reset** to adjust the time stamp back to the original times.
- 5 Click **Apply**.

The log browser displays the log entries with the adjusted times.

## Export Logs from the Log Browser

You can export log files using the Log Browser.

#### Procedure

- 1 Navigate to the Log Browser and select an object to browse.
- 2 Select **Action > Export**.
- 3 Select the type of file that you want to export.

Selecting the log file bundle downloads a potentially large file.

- 4 Click **Export**.  
After a few seconds a new web browser displays.
- 5 Enter the location where you want to save the file.

The log file is downloaded to your local machine and you can close the new browser window.

## Compare Log Files

You can open multiple windows in the Log Browser to compare log files.

### Procedure

- 1 Navigate to the Log Browser and retrieve a log file from an object.
- 2 Select **Actions > New Browser Window** to open a window in the Log Browser.
- 3 In the Log Browser window, retrieve another log file to view.

You can perform the same actions with the log file opened in the new window as you can with the original Log Browser window.

## Manage Logs Using the Log Browser

From the Log Browser, you can update, remove, and see a list of available log file bundles.

To manage log file bundles, you must access the Log Browser from the vSphere Web Client home.

### Procedure

- 1 From the vSphere Web Client home, click **Log Browser**.
- 2 Click the **Manage** tab.
- 3 Select an object's log file in the list of sources.
- 4 (Optional) Click **Remove** to delete the log file bundle.

Deleting the log file bundle reclaims disk space on the vSphere Web Client server.

All log files generated from that bundle are deleted.

- 5 (Optional) Click **Update** to update the list of log file bundles.

You can view the log bundles created by other vSphere Web Client sessions.

The log does not appear in the retrievable objects list.

## Browse Log Files from Different Objects

You can browse multiple log files coming from different objects within the Log Browser at the same time. It helps in comparing log files simultaneously.

### Procedure

- 1 From the vSphere Web Client home, click **Log Browser**.
- 2 Click the **View** tab.
- 3 To view its logs, select an object (ESXi host or vCenter Server)
- 4 Open a new browser window by selecting **Actions > New Browser Window** and select another object to view its logs.