Dell EMC VxFlex Ready Node

Version 2.x

AMS User Guide

P/N 302-003-031

REV 06

DELL EMC

Copyright © 2016-2018 Dell Inc. or its subsidiaries. All rights reserved.

Published June 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
1-508-435-1000 In North America 1-866-464-7381
www.DellEMC.com

CONTENTS

Figures 11

Tables 13

Preface 15

Part 1 Introduction 17

Chapter 1 Introduction to VxFlex Ready Node 19
What is VxFlex OS? 20
System requirements 21
Server requirements 21
Other requirements 22
Boot drive 22
Switch requirements 23
IP addresses 23
New and changed features 23
New features for VxFlex OS 24
New features for VxFlex Ready Node 24
Product limits 24

Chapter 2 Architecture 27
VxFlex OS architecture overview 28
System 28
Hardware 28
Software 28
The MDM cluster 29
Storage definitions 31
Protection Domains 31
Storage Pools 31
Fault Sets 33
Naming 34
Protection and load balancing 34
Rebuild 34
Rebuild throttling 34
Rebalance 35
Rebalance throttling 36
Checksum protection 36
ESX vStorage APIs for Array Integration (VAAI) 37
Caching 38
Networking 40
IP Addresses 40
VLAN support 42
Networking and connectivity - 13G servers 43
Networking and connectivity - 14G servers 50
Virtual IP Address 60
Monitoring of SDC and SDS connections 61

Snapshots.....61
V-Trees.....62
Other functions.....63
Implementing VxFlex OS.....64
Implementing VxFlex OS in an ESXi-based system.....64
Implementing VxFlex OS over a Linux bare metal system.....66
SAN virtualization layer.....67
Compute cluster configuration and workload deployment.....68
Maintenance.....69
Maintaining the physical layer.....69
Instant maintenance mode.....70
Maintaining the virtualization layer.....70
Management tools.....71
Firmware patching.....71

Part 2 Getting Started 75

Chapter 3 User Management 77
Overview.....78
User roles.....78
Authenticating with the management server.....79
Logging in.....80
Logging out.....81
Setting the User Authentication Method.....82
Adding and modifying local users.....83
Adding users.....83
Deleting users.....84
Modifying user credentials.....84
Displaying system users and their roles.....85
Displaying information about a specific user.....86
Resetting user passwords.....86
Changing user passwords.....87
Disabling the Default Admin User.....88

Chapter 4 Creating and Mapping Volumes 91
Creating and mapping volumes overview.....92
Creating volumes.....92
Adding volumes.....92
Mapping a volume to an SDC.....94
Mapping volumes.....94

Part 3 Managing and Monitoring 97

Chapter 5 Managing System Objects 99
CLI basics.....100
Using SCLI in non-secure mode.....101
Syntax.....101
Getting help with the CLI.....102
Managing the VxFlex OS system.....103
Extending an existing VxFlex OS system.....103
Managing the MDM cluster.....103

Replacing, or updating an IP address on a member of the MDM cluster....103
Transfer the MDM cluster component using the AMS....110
Removing a standby MDM with the VxFlex OS GUI....112
Configure virtual IP addresses....112
Renaming an MDM Cluster Node....113
Managing the SDSs and cache....113
Modifying an SDS port during I/O....114
Managing Read Flash cache....115
Managing read RAM cache....116
SDC operations....116
Updating the SDC driver with IP changes....116
Detecting new volumes....117
Query volumes using drv_cfg....118
Query tgt objects using drv_cfg....118
Query GUID using drv_cfg....119
Query MDMs using drv_cfg....120
Loading a configuration file using drv_cfg....120
Adding an MDM using drv_cfg....121
Modifying an MDM IP address using drv_cfg....122
Permanent Device Loss state....123
Managing ESX servers....124
Modifying parameters on ESX servers....124
Checking the SDC state on ESX servers....126

Chapter 6 Security Management 129
Sign and update AMS security certificates....130
Signing AMS security certificates using CLI commands....131
Modifying AMS-MDM connection security....132
Modifying AMS-MDM connection security....134
Back up SSH public keys using the VxFlex OS GUI....136
Backing up SSH public keys using CLI commands....137
Changing the emergency password....138
Setting credentials for existing Monitor users....138
Changing login credentials for existing Management users....139
Login banner overview....140
Setting up a login banner using the CLI....140
Enabling or disabling preemptive acceptance of the login banner....141
Activating preemptive acceptance of the login banner....141

Chapter 7 Opening the VxFlex OS GUI and Logging In 143
Log in to the VxFlex OS GUI....144
Connection and disconnection information....144

Chapter 8 VxFlex OS GUI Features 145
VxFlex OS GUI overview....146
VxFlex OS GUI conventions....147
Alerts indicators....147
Color codes....147
Hardware view....148
Dashboard view....150
Navigation tree....151
System Settings and Logout menus....151

Dimmer 152
Widget 152
Dashboard tiles 152
Management (MDM) cluster status 155
Frontend views 157
Volumes 158
SDCs 159
Snapshots 159
Adding Frontend objects to the filter 159
Compute 160
Backend view 161
Command menu (3) 163
Filter (1) 163
Table 163
Alerts view 167
Property Sheets 168

Chapter 9 Monitoring the System using the VxFlex OS GUI 173
Viewing object properties in Backend and Frontend views 174
Viewing detailed hardware properties 174
Identifying hardware errors 174
Identifying the IP role of a node 174
Identifying the SDS running on a node 174
Identifying the SSD used for caching 175
MTU values 176
Monitoring compute and network environments 176
Viewing licensing information 177
Verifying your connection to the Management cluster 178
Collecting logs 178

Chapter 10 Configuring the System using the VxFlex OS GUI 179
Configuring capacity 180
Adding nodes 180
Remove the node from the system 180
Adding devices 182
Removing devices 183
Turning on identification LEDs 187
Adding, removing, and activating and inactivating capacity 187
Activating devices 190
Clearing device errors 190
Setting device capacity limits 190
Modifying cache settings 191
Enabling and disabling caching at Storage Pool level 191
Designating an SSD as a caching device 192
Stopping caching on a specific device 192
Removing an SSD device from the caching pool 192
Configuring Read Flash Cache settings per Protection Domain 193
Changing Read RAM Cache volume settings 193
Clearing device errors 194
Configuring Read RAM Cache (advanced, Backend) 194
Configuring volumes, volume trees, SDCs, and snapshots 196
Adding volumes 196
Restricted SDC mode 198
Mapping volumes 199

Unmapping volumes..... 200
Removing volumes.....200
Creating a datastore.....201
Deleting a datastore..... 202
Creating volume snapshots.....202
Removing snapshots.....203
Increasing a volume's size.....204
Setting volume bandwidth and IOPS limits..... 204
Entering and exiting SDS Instant Maintenance Mode.....205
Cluster operations..... 206
Adding a standby to the MDM cluster.....207
Removing a standby from the MDM cluster.....207
Extending an MDM cluster..... 208
Switching from 5-node cluster mode to 3-node cluster mode..... 209
Replacing a Master MDM..... 209
Transfer the MDM cluster component using the AMS..... 210
Configuring Oscillating Failure counters..... 212
Configuring Oscillating Failure counter parameters..... 212
Resetting Oscillating Failure counters..... 213
Viewing Oscillating Failure counters..... 214
Applying Performance Profiles to system components..... 215
Configuring I/O priorities and bandwidth use (advanced)..... 216
Application IOPS and bandwidth (advanced)..... 216
System IOPS and bandwidth (advanced).....216
Enable and disable Rebuild/Rebalance (advanced)..... 217
Using the background device scanner..... 218
Enabling and disabling the background device scanner..... 218
Resetting the background device scanner counters..... 219
Modifying Checksum protection mode.....219
Renaming objects..... 220
Provision compute, storage and network environments, and deploy workloads 220
Configure VMkernel IP address pools.....221
Provision an environment in simple mode..... 222
Provision an environment in advanced mode..... 224
Deploy a workload.....228
Expand environments..... 230
Delete environments..... 231
Remove a node from a compute cluster..... 231
Expand a workload..... 231
Delete a workload..... 232
Compute, storage and network environment default values..... 232
Approving pending security certificates..... 235
Customizing system preferences.....235
Modifying system settings.....238
Management and Data Network IP addresses..... 238
Default DNS and NTP IP addresses..... 239
ESRS..... 239
SNMP..... 241
Install the license - using the VxFlex OS GUI..... 242
Setting the vCenter..... 242
Discovery.....243
Changing login credentials for existing Management users.....243

Part 4 Reference 245

Chapter 11 Common Tasks 247
Associating VxFlex OS volumes with physical disks.....248
Volume information - Linux.....248
Port usage and changing default ports.....249
Upgrading Java.....250
Move ESXi hosts to a new vCenter and reconfigure AMS.....250
Best practices for using VMware vMotion.....251
Disable VMware Admission Control.....252
Update the hardware using the iDRAC virtual console.....253
Update the SATADOM firmware (13G servers only).....254
AMS management takeover of non-managed nodes.....255
AMS management takeover compatibility.....255
AMS management takeover of non-managed nodes.....256

Chapter 12 System events 263
System events overview.....264
Event format.....264
Viewing events.....265
Viewing events locally.....265
Event list.....268
Authentication.....268
CLI commands.....269
License and installation.....270
MDM.....272
SDC.....274
SDS.....275
Rebuild.....278

Chapter 13 SNMP Trap Support 281
General.....282
Supported alerts and event numbering conventions.....282
CLASS/TYPE:.....282
VxFlex OS Alerts in SNMP, VxFlex OS GUI, REST, and ESRS.....283
VxFlex Ready Node–specific Alerts in SNMP, VxFlex OS GUI, REST, and ESRS.....300
Configure SNMP properties after deployment.....308
Create a lockbox.....309
Configuring SNMP using the VxFlex OS GUI.....310
VxFlex OS.mib file.....310

Chapter 14 Server Factory Reset and Power Operations 313
Perform a factory reset.....314
Set up the BMC (iDRAC) IP address and BIOS.....314
Verify the status of the system hardware, storage controller, and disks - 13G servers.....316
Verify the status of the system hardware, storage controller, and disks - 14G servers.....317
Open the KVM console - 13G servers.....318
Open the KVM console - 14G servers.....319
Map the ISO file on a VxFlex Ready Node server.....320
Map the ISO file on a VxFlex Ready Node 14G server.....321
KVM activities.....322
Additional tasks using the KVM console - 13G server.....322

Additional tasks using the KVM console for 14G servers..... 323

Chapter 15 DTK - Hardware Update Bootable ISO 325

Dell OpenManage DRAC Tools (RACADM).....326
Update the hardware using remote RACADM.....326
Update the SATADOM firmware (13G servers only).....328
Recommended BIOS and firmware settings.....329
VxFlex OS ID module..... 329
BIOS and firmware..... 329
Configuration settings..... 330
Troubleshooting the Hardware ISO..... 331
Troubleshoot general iDRAC failures..... 331
iDRAC virtual console issues.....332
iDRAC virtual media issues.....332
DTK A08: BIOS set commands fail.....333
Rebranding operation fails..... 333
Check the logs for error messages..... 333
Additional resources..... 334
VxFlex OS resources.....334
Dell Lifecycle Controller (LC)..... 334
Dell OpenManage Deployment Toolkit (DTK)..... 335

Glossary 337

FIGURES

1 5-node MDM cluster ... 31
2 Protection Domains and Storage Pools ... 32
3 Protection Domains, Storage Pools, and Fault Sets ... 33
4 Single-node R730xd (2U1N) server (storage only) typical connection diagram ... 44
5 Single-node R730xd (2U1N) server (converged) typical connection diagram ... 44
6 Single-node R630 (1U1N) server typical connection diagram ... 45
7 Correlating physical slot and Ethernet network names ... 50
8 R640 and R740xd PCI slots, integrated NICs and BMC port locations ... 52
9 Correlating physical slot and Ethernet network names ... 60
10 Snapshot operations ... 62
11 V-Tree diagram ... 63
12 VxFlex OS implementation on ESX ... 65
13 Add Volume window ... 94
14 Map Volumes window after mapping is complete ... 95
15 Target server ... 111
16 Source server, now a standby MDM ... 112
17 Target server, now an MDM ... 112
18 System Settings window, Technical Users pane, Monitor User ... 139
19 Alerts indicators ... 147
20 Hardware view ... 149
21 Dashboard controls ... 150
22 Widget, showing Return to Full-Screen button ... 152
23 Dashboard tiles ... 153
24 Frontend > Volumes view ... 158
25 Frontend > SDCs view ... 159
26 Frontend > filtered Snapshots view ... 159
27 Frontend > Compute view ... 160
28 Backend view ... 162
29 Backend filter ... 163
30 Alerts view ... 167
31 Multiple floating Property Sheets ... 171
32 Example of a Property Sheet for an SDS ... 172
33 A selected device (1) and its corresponding Property Sheet (2) ... 175
34 Examples of drives: type (1, 2), storage (3), and caching (4) ... 176
35 Faulty server in BackendStorage view ... 181
36 Set Volume Use Read RAM Cache window ... 194
37 Read RAM Cache configuration at Storage Pool level ... 195
38 Read RAM Cache configuration at SDS level ... 195
39 Add Volume window ... 197
40 Target server ... 211
41 Source server, now a standby MDM ... 212
42 Target server, now an MDM ... 212
43 Clusters shown in VxFlex OS GUI and in vCenter ... 224
44 Clusters shown in VxFlex OS GUI and in vCenter ... 228
45 Workloads shown in VxFlex OS GUI and in vCenter ... 230
46 User Preferences window ... 236
47 Network ports ... 252

TABLES

1 Server requirements....21
2 Product limits.... 24
3 MDM cluster modes.... 30
4 Caching modes.... 38
5 Caching support matrix.... 40
6 VxFlex OS management IP network....40
7 VxFlex OS data IP network (continued).... 41
8 R730xd single-node server configuration.... 45
9 R630 single-node server configuration.... 46
10 R630 single-node server configuration.... 47
11 R730xd single-node server configuration....47
12 R730xd single-node server configuration (Storage Only).... 48
13 R740xd server configuration for 10/25 GbE, 32xSAS, SFP+/SFP28....53
14 R740xd server configuration for 10/25GbE, 12xNVMe, SFP+/SFP28.... 54
15 R740xd server configuration for 25GbE, 3xGPU (M60)/ 2xGPU (M10), SFP28.... 55
16 R740xd server configuration for 10GbE, 3xGPU (M60)/ 2xGPU (M10), SFP+....55
17 R640 server configuration for 10/25GbE, 1CPU, SFP+/SFP28....56
18 R640 server configuration for 10/25GbE, 2CPU, SFP+/SFP28.... 56
19 R640 single-node server configuration.... 57
20 R740xd single-node server configuration for 10/25GbE, 2CPU, SFP+/SFP28....58
21 13G PowerEdge Servers....71
22 14G PowerEdge Servers.... 72
23 Local and LDAP user roles and permissions.... 78
24 VxFlex OS GUI task overview....146
25 VxFlex OS GUIcolor codes....147
26 Alert symbols and color codes....148
27 Dashboard tiles.... 153
28 Management node icons with normal operational status (green).... 156
29 Management node status indications and color codes....157
30 Table view options....164
31 User Preferences....236
32 Default ports.... 249
33 VxFlex OS Alerts in SNMP, VxFlex OS GUI, REST, and ESRS.... 283
34 VxFlex Ready Node–specific Alerts in SNMP, VxFlex OS GUI, REST, and ESRS..... 300
35 Hardware ISO configuration settings....330

Preface

As part of an effort to improve its product lines, Dell EMC periodically releases revisions of its software and hardware. Therefore, some functions described in this document might not be supported by all versions of the software or hardware currently in use. The product release notes provide the most up-to-date information on product features.

Contact your Dell EMC technical support professional if a product does not function properly or does not function as described in this document.

Note

This document was accurate at publication time. Go to Dell EMC Online Support (https://support.emc.com) to ensure that you are using the latest version of this document.

Previous versions of Dell EMC VxFlex OS were marketed under the name Dell EMC ScaleIO.

Similarly, previous versions of Dell EMC VxFlex Ready Node were marketed under the name Dell EMC ScaleIO Ready Node.

References to the old names in the product, documentation, or software, etc. will change over time.

Note

Software and technical aspects apply equally, regardless of the branding of the product.

Related documentation

The release notes for your version includes the latest information for your product.

The following Dell EMC publication sets provide information about your VxFlex OS or VxFlex Ready Node product:

- VxFlex OS software (downloadable as VxFlex OS Software <version> Documentation set)
- VxFlex Ready Node with AMS (downloadable as VxFlex Ready Node with AMS Documentation set)
- VxFlex Ready Node no AMS (downloadable as VxFlex Ready Node no AMS Documentation set)
- VxRack Node 100 Series (downloadable as VxRack Node 100 Series Documentation set)

You can download the release notes, the document sets, and other related documentation from Dell EMC Online Support.

Typographical conventions

Dell EMC uses the following type style conventions in this document:

Bold	Used for names of interface elements, such as names of windows, dialog boxes, buttons, fields, tab names, key names, and menu paths (what the user specifically selects or clicks)

Italic	Used for full titles of publications referenced in text
`Monospace`	Used for: • System code • System output, such as an error message or script • Pathnames, filenames, prompts, and syntax • Commands and options
Monospace italic	Used for variables
`Monospace bold`	Used for user input
[ ]	Square brackets enclose optional values
|	Vertical bar indicates alternate selections - the bar means “or”
{ }	Braces enclose content that the user must specify, such as x or y or z
...	Ellipses indicate nonessential information omitted from the example

Where to get help

Dell EMC support, product, and licensing information can be obtained as follows:

Product information

For documentation, release notes, software updates, or information about Dell EMC products, go to Dell EMC Online Support at https://support.emc.com.

Technical support

Go to Dell EMC Online Support and click Service Center. You will see several options for contacting Dell EMC Technical Support. Note that to open a service request, you must have a valid support agreement. Contact your Dell EMC sales representative for details about obtaining a valid support agreement or with questions about your account.

Your comments

Your suggestions will help us continue to improve the accuracy, organization, and overall quality of the user publications. Send your opinions of this document to techpubcomments@emc.com.

PART 1

Introduction

This section describes an overview of the benefits and system requirements of VxFlex OS.

Chapters include:

Chapter 1, "Introduction to VxFlex Ready Node"

Chapter 2, "Architecture"

CHAPTER 1

Introduction to VxFlex Ready Node

This section introduces VxFlex Ready Node.

- What is VxFlex OS? 20
- System requirements 21
- New and changed features 23
- Product limits 24

What is VxFlex OS?

VxFlex OS

VxFlex OS, based on the VxFlex OS software platform, is a software-only solution that uses existing servers' local disks and LAN to create a virtual SAN that has all the benefits of external storage—but at a fraction of cost and complexity. VxFlex OS utilizes the existing local storage devices and turns them into shared block storage. For many workloads, VxFlex OS storage is comparable to, or better than external shared block storage.

The lightweight VxFlex OS software components are installed on the application servers and communicate via a standard LAN to handle the application I/O requests sent to VxFlex OS block volumes. An extremely efficient decentralized block I/O flow, combined with a distributed, sliced volume layout, results in a massively parallel I/O system that can scale up to thousands of nodes.

VxFlex OS is designed and implemented with enterprise-grade resilience. Furthermore, the software features an efficient distributed self-healing process that overcomes media and server failures, without requiring administrator involvement.

Dynamic and elastic, VxFlex OS enables administrators to add or remove servers and capacity on-the-fly. The software immediately responds to the changes, rebalancing the storage distribution and achieving a layout that optimally suits the new configuration.

Because VxFlex OS is hardware agnostic, the software works efficiently with various types of disks, including: magnetic (HDD) and solid-state disks (SSD), flash PCI Express (PCIe) cards, networks, and hosts.

VxFlex OS can easily be installed in an existing infrastructure as well as in green field configurations.

VxFlex Ready Node

VxFlex Ready Node is the combination of VxFlex OS software-defined block storage and Dell PowerEdge® servers, optimized to run VxFlex OS, enabling customers to quickly deploy a fully architected, software-defined, scale out server SAN.

Any hypervisor can run on VxFlex Ready Node servers as an application consuming VxFlex OS volumes.

In an AMS-based solution, a limited number of ESXi and RHEL operating systems are currently supported. For more information, see the operating system support tables.

The solution is managed by the AMS (Automated Management Services), which enables a simple or customized deployment process from bare metal, no IP state, to a fully-configured system: IP address assignment, VxFlex OS deployment and configuration, and vCenter configuration.

Note

In the documentation set, there are references that are specific to either ESXi or RHEL operating systems, but not to both (for example, vCenter). These differences are not marked in most places.

The AMS solution enables the following:

- VxFlex OS GUI client
- CLI client (SCLI and AMS CLI)
- REST

The AMS enables reporting via ESRS (EMC proprietary) and SNMP.

The minimum system size is three nodes. The hardware consists of single nodes with various configurations of SAS HDD, SSD, and NVMe SSD devices. Some of the configurations include hardware read and write-back caching ability, running on a dedicated SSD.

VxFlex Ready Node gives administrators a fully-converged solution that is easy to deploy, configure, and upgrade.

System requirements

This section lists the requirements for system components.

This section is specific to VxFlex OS software deployments.

For VxFlex Ready Node systems, refer to your product's Hardware Configuration and Operating System Installation Guide.

Server requirements

Table 1 Server requirements

Component	Requirement
Management server	• Processor - 2 vCPUs • Memory- 3 GB • Storage - 80 GB of free storage capacity
VxFlex OS GUI requirements	• Java v1.8 (64-bit), build 149 or earlier. You can Download previous versions from this link: http://www.oracle.com/technetwork/java/javase/downloads/java-archive-javase8-2177648.html • Supported operating systems: ▪ Windows 7 ▪ Windows 10 ▪ Windows 2012 R2 ▪ Windows 2016
Node operating system	• ESXi 6.0 P6 and 6.5 P2 • RHEL 7.3 and 7.4 (patched with Meltdown and Specter fixes)
AMS server	• Operating systems (all 64-bit): ▪ Windows 2012 R2, 2016. Server Core editions are not supported. **Note** You must install the Microsoft Visual C++ 2010 Redistributable Package (x64). This package should be part of the operating system image.

Table 1 Server requirements (continued)

Component	Requirement
	▪ RHEL 7.2, 7.3, 7.4 ▪ SLES 12.2 **Note** For Linux servers, ensure that the following packages are installed: – binutils – libnuma-devel – libnuma1 • Java v1.8 (64-bit), build 149 or earlier. You can Download previous versions from this link: http://www.oracle.com/technetwork/java/javase/downloads/java-archive-javase8-2177648.html
Host image recommendations	• For the latest releases of ESXi 6.0/6.5 and RHEL 7.3/7.4, no optimization is required.
KVM browser requirements	• Firefox 30 or later • Internet Explorer 8 or later • Java 1.8

For the latest operating system support matrix, see the ESSM (https://elabnavigator.emc.com).

Other requirements

VxFlex OS requires that you use a minimum of three SDS servers, with a combined free capacity of at least 300 GB. These minimum values are true per system and per Storage Pool.

All VxFlex Ready Node nodes have the minimum disk size as part of the node specification.

The AMS server must be a separate server that resides outside of the VxFlex OS system.

Supported hardware configurations
The supported hardware configurations are listed in the EMC Simple Support Matrix for your software version. Download this file from https://elabnavigator.emc.com/eln/elnhome. Click **Simple Support Matrices** > **Storage** and select the desired product and version.

Boot drive

All configurations have one of the following boot drive configurations:

- 64 GB SATADOM

- 2 120/240 GB M.2 devices (RAID 1)

Switch requirements

One 10/25/100 GbE switch is required. Multiple switches may be used for high availability and load balancing purposes. Separate (one or more) 1 GbE switches may be used for the VxFlex OS management network.

If the node doesn't have a 1 GbE management port, the 1 GbE switch is used for the iDRAC port and the leftmost 10\25 GbE port on the node is used for management and will be connected to a 10/25 GbE switch.

Before you commence connecting the VxFlex OS system to the switches, ensure that they are installed and configured. Do not connect the switches to your organization's network until all installation and configuration activities are complete.

The following use cases are supported:

- 10 GbE
  Each node has at least four 10 GbE ports. Two are used by VxFlex OS for internal network communication. The other two ports are used for VMware high availability, DRS, and application network traffic.
- 25/100 GbE
  - Each node has at least four 25 GbE ports. Two are used by VxFlex OS for internal network communication. The other two ports are used for VMware high availability, DRS, and application network traffic.
  - Each node has two 25 GbE ports, both used for VxFlex OS internal network communication and application traffic.

The switches must have sufficient available network ports to accommodate the following:

Data network 10/25/100 GbE switches:

- Two 10/25/100 GbE ports per node, per switch

Management network switches:

- One 1 GbE and one BMC (iDRAC) port per node.
- If the node does not have a 1 GbE port, minimally a BMC (iDRAC) port is needed.

IP addresses

IP Addresses on page 40 explains the required network configurations for VxFlex Ready Node.

vCenter support

vCenter 6.0 and 6.5 are supported.

New and changed features

New features for VxFlex OS

Learn about the new features introduced in the VxFlex OS 2.6-<build> software.

Security fixes
Spectre and Meltdown issues have been addressed in this version. For more information, refer to the *VxFlex OS Security Configuration Guide*.

New features for VxFlex Ready Node

Learn about the new features introduced in VxFlex Ready Node 2.6-<build>.

Security fixes
Spectre and Meltdown issues have been addressed in this version.

- The Automated Management Service (AMS) integrates the fixes as part of the initial deployment for new systems. For upgrades, the *VxFlex Ready Node Upgrade Guide* includes instructions for implementing the security fixes, as part of the upgrade process.
- In AMS managed systems, on Linux-based servers, the administrator can enable or disable patches on the server. For more information, refer to the *VxFlex Ready Node Security Configuration Guide*.

Product limits

The following table lists product capabilities:

Table 2 Product limits

Item VxFlex OS	Limit
System raw capacity	300 GB—16 PB
Device size	100 GB—8 TB
Minimum Storage Pool capacity	300 GB
Volume size	8 GB—1 PB
Maximum number of volumes/snapshots in system	32,768[a]
Maximum number of volumes/snapshots in Protection Domain	32,768
Maximum number of volumes + snapshots in single VTree	32
Maximum capacity per SDS	96 TB
Maximum Nodes per system (includes SDSs and SDCs)	1024
SDSs per Protection Domain	128[a]
Maximum devices (disks) per SDS server	64[b]
Maximum devices (disks) per Storage Pool	300[a]
Minimum devices (disks) per Storage Pool	3, on different SDSs

Table 2 Product limits (continued)

Item VxFlex OS	Limit
Maximum volumes that can be mapped to a single SDC	256
Maximum Protection Domains per system	256
Maximum Storage Pools	1024
Maximum Storage Pools per Protection Domain	64
Maximum Fault Sets per Protection Domain	64
Maximum IP addresses per server (MDM and SDS)	8
RAM Cache	128 MB—300 GB

a. If more are needed, contact Customer Support.
b. On VMware servers, the maximum devices per SDS is 59.

CHAPTER 2

Architecture

The following topics describe the VxFlex Ready Node architecture.

- VxFlex OS architecture overview 28
- System 28
- The MDM cluster 29
- Storage definitions 31
- Protection and load balancing 34
- ESX vStorage APIs for Array Integration (VAAI) 37
- Caching 38
- Networking 40
- Virtual IP Address 60
- Monitoring of SDC and SDS connections 61
- Snapshots 61
- Other functions 63
- Implementing VxFlex OS 64
- SAN virtualization layer 67
- Compute cluster configuration and workload deployment 68
- Maintenance 69
- Management tools 71
- Firmware patching 71

VxFlex OS architecture overview

This chapter describes the VxFlex OS architecture overview.

VxFlex OS is a software-only solution. VxFlex OS components are lightweight, highly available software components, installed on new or existing servers alongside your production applications (hypervisors, databases, web applications, etc.). The system can be installed directly on the servers, or over a virtual server system (hypervisor or virtual machines).

System

VxFlex OS is based on a hardware and a software component.

Hardware

The VxFlex Ready Node system includes dedicated server nodes, built for capacity, performance, or compute, per the specific configuration.

The system refers to the following hardware components:

- *Nodes*
  Nodes are the basic computer unit used to install and run the operating system and VxFlex OS. They can be the same servers used for the applications (server convergence), or in a 2-layer architecture.

  We use these terms to describe the hardware in the system:

  - Node
    A server, in one of the following forms:
    - 2U1N node—the node takes the entire space of the chassis. This is like a standard rack-mounted server. The chassis is 2U high.
    - 1U1N node—the node takes the entire space of the chassis. This is like a standard rack-mounted server. The chassis is 1U high.
- *Storage Media*
  The storage media is supplied as part of the hardware offering.

Software

The VxFlex OS virtual SAN consists of the following software components:

- Meta Data Manager (MDM)
  Configures and monitors VxFlex OS. The MDM can be configured in redundant cluster mode, with three members on three servers or five members on five servers, or in single mode on a single server.

  NOTICE

  It is not recommended to use single mode in production systems, except in temporary situations. The MDMs contains all the metadata required for system operation. single mode has no protection, and exposes the system to a single point of failure.

- Storage Data Server (SDS)
  Manages the capacity of a single server and acts as a back-end for data access. The SDS is installed on all servers contributing storage devices to VxFlex OS. These devices are accessed through the SDS.

- Storage Data Client (SDC)
  A lightweight device driver that exposes VxFlex OS volumes as block devices to the application that resides on the same server on which the SDC is installed.
- Automated Management Services (AMS)
  A dedicated management server, residing outside the VxFlex Ready Node converged system, that serves VxFlex OS GUI, CLI, and REST clients.

Depending on the desired configuration (described later), the software components are installed on the server and give rise to a virtual SAN layer exposed to the applications that reside on the servers.

The MDM cluster

The MDM serves as the monitoring and configuration agent of VxFlex OS. The MDM is mainly used for management which consists of migration, rebuilds, and all system-related functions. No I/O run through the MDM.

To support high availability, three or more instances of MDM run on different servers. In a multi-MDM environment, one MDM is given the Master role, and the others act as Slave or Tie Breaker MDMs.

The MDM cluster comprises a combination of Master MDM, Slave MDMs, and Tie Breaker MDMs.

The following terms are relevant to the MDM, the building blocks of the MDM cluster:

- MDM

  Any server with the MDM package installed on it. An MDM can be given a Manager or a Tie Breaker (default) role, during installation. MDMs have a unique MDM ID, and can be given unique names.

  Before the MDM can be part of the cluster, it must be promoted to a Standby MDM.
- Standby MDM and Tie Breaker

  An MDM and a Tie Breaker can be added to a system as a standby. Once added, the standby MDM or Tie Breaker is attached, or locked, to that specific system.

  A standby MDM can be called on to assume the position of a Manager MDM or Tie Breaker MDM when it is promoted to be a cluster member.
- Manager MDM

  An MDM that can act as a Master or a Slave in the cluster. Manager MDMs have a unique system ID, and can be given unique names. A manager can be a standby or a member of the cluster.

  In VxFlex OS documentation, “MDM” refers to a manager, unless specified otherwise.
- Tie Breaker MDM

  An MDM whose sole role is to help determine which MDM is the master. A Tie Breaker can be a standby or a member of the cluster. A Tie Breaker MDM is not a manager.

  In a 3-node cluster, there is one TB; in a 5-node cluster, there are two TBs. This ensures that there are always an odd number of MDMs in a cluster, which guarantees that there is always a majority in electing the master.

The following terms are relevant to the MDM cluster, specifically:

- Master MDM (used to be called Primary MDM)

The MDM in the cluster that controls the SDSs and SDCs. The Master MDM contains and updates the MDM repository, the database that stores the SDS configuration, and how data is distributed between the SDSs in the system. This repository is constantly replicated to the Slave MDMs, so they can take over with no delay.

Every MDM cluster has one Master MDM.

- Slave MDM (used to be called Secondary MDM)

  An MDM in the cluster that is ready to take over the Master MDM role if ever necessary.

  In a 3-node cluster, there is one Slave MDM, thus allowing for a single point of failure. In a 5-node cluster, there are two Slave MDMs, thus allowing for two points of failure. This increased resiliency is a major benefit to enabling the 5-node cluster.

- Replica

  An MDM that contains a replica of the MDM repository. This includes the Master MDM and any Slave MDMs in the MDM cluster.

The following table describes the available cluster modes:

Table 3 MDM cluster modes

Cluster mode	Members	Description
3-node (default)	• Master MDM • Slave MDM • Tie Breaker	3-node cluster has two copies of the repository, thus can withstand one MDM cluster failure.
5-node	• Master MDM • Two Slave MDM • Two Tie Breaker	5-node cluster has three copies of the repository, thus can withstand two MDM cluster failure.
Single-node	• Master MDM	Single-node cluster has only one copy of the repository, thus it cannot withstand failure. It is not recommended to use Single Mode in production systems, except in temporary situations.

In addition to the cluster members, you can prepare standby Managers and Tie Breaker nodes, for a total of thirteen cluster and standby MDMs.

The MDM cluster IP address limit is 16 IP addresses, which includes all cluster members (Master, Slave, Standby Master, and Standby Slaves).

The following figure illustrates a 5-node MDM cluster:

Figure 1 5-node MDM cluster

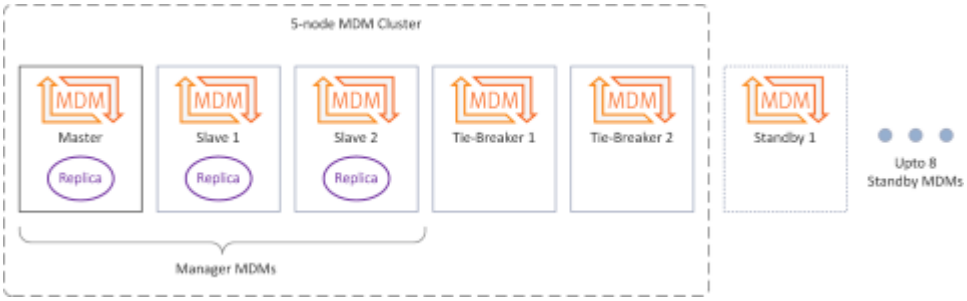

All members of the MDM cluster have the same MDM package installed on them.

Before a server makes its way into the MDM cluster, it must follow the following path:

1. Install the MDM package on the server.
   During the installation, you determine if the server will be a Manager or a Tie Breaker (default).
2. Promote the server to Standby status, either as a Manager or as a Tie Breaker.
3. Add the standby server to the MDM cluster. A Manager, once entered into the cluster can take on the Master or Slave state.

MDM cluster creation is done automatically when deploying a system with any of the automated deployment tools.

Storage definitions

When configuring VxFlex OS, you should take the following concepts into account: Protection Domains, Storage Pools, and Fault Sets. Together, these elements link the physical layer with the virtualization layer.

Protection Domains

A Protection Domain is a logical entity that contains a group of SDSs that provide backup for each other. Each SDS belongs to one (and only one) Protection Domain. Thus, by definition, each Protection Domain is a unique set of SDSs. In Figure 2 there are three Protection Domains. The one in the middle (fully depicted) consists of seven SDSs, each with two storage devices.

The maximum recommended number of nodes in a Protection Domain is 100. This enables the following:

- optimal performance
- reduction of theoretical *mean time between failure* issues
- ability to sustain multiple failures in different Protection Domains

You can add Protection Domains during installation. In addition, you can modify Protection Domains post-installation with all the management clients .

Storage Pools

Storage Pools allow the generation of different storage tiers in VxFlex OS. A Storage Pool is a set of physical storage devices in a Protection Domain. Each storage device belongs to one (and only one) Storage Pool. In Figure 2, there are two Storage Pools depicted.

When a volume is configured over the virtualization layer (see "SAN virtualization layer"), it is distributed over all devices residing in the same Storage Pool. Each

volume block has two copies located on two different SDSs. This allows the system to maintain data availably following a single-point failure. The data will still be available following multiple failures, as long as each failure took place in a different Storage Pool.

Figure 2 Protection Domains and Storage Pools

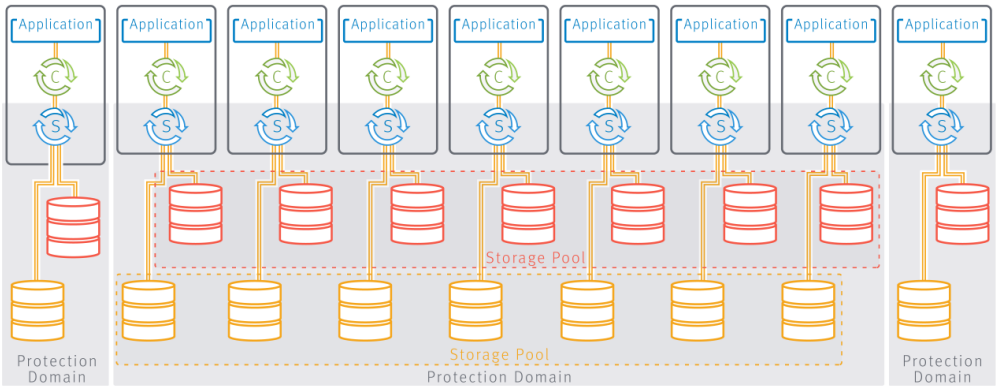

All devices in a Storage Pool must be of the same device media type.

For example, consider Figure 2. If all SDSs in a Protection Domain have two physical drives associated with them—one HDD, and the other SSD— you should define two Storage Pools:

- Magnetic Storage Pool

  Consists of all HDDs in the Protection Domain

- High performance Storage Pool

  Consists of all SSDs used for storage purposes in the Protection Domain

---

Note

VxFlex OS might not perform optimally if there are large differences between the sizes of the devices in the same Storage Pool. For example, if one device has a much larger capacity than the rest of the devices, performance may be affected. After adding devices, you can define how much of the device capacity is available to VxFlex OS by using the SCLI `modify_sds_device_capacity` command.

Each Storage Pool can work in one of the following modes:

- Zero-padding enabled

  Ensures that every read from an area previously not written to returns zeros. Some applications might depend on this behavior. Furthermore, zero padding ensures that reading from a volume will not return information that was previously deleted from the volume.

  This behavior incurs some performance overhead on the first write to every area of the volume.

- Zero-padding disabled (default)

  A read from an area previously not written to will return unknown content. This content might change on subsequent reads.

Zero padding must be enabled if you plan to use any other application that assumes that when reading from areas not written to before, the storage will return zeros or consistent data.

Note

The zero padding policy cannot be changed after the addition of the first device to a specific Storage Pool.

You can add Storage Pools during installation. In addition, you can modify Storage Pools post-installation with most of the management clients.

Fault Sets

A Fault Set is a logical entity that contains a group of SDSs within a Protection Domain, that have a higher chance of going offline together, such as a group of SDSs that are all powered in the same rack. By grouping them into a Fault Set, you are configuring VxFlex OS to mirror the data for all devices in this Fault Set. The mirroring should take place on SDSs that are outside of this Fault Set.

When defining Fault Sets, we refer to the term fault units, where a fault unit can be either a Fault Set, or an SDS not associated with a Fault Set (you may think of it as a Fault Set of a single SDS).

There must be enough capacity within at least three fault units to enable mirroring.

If Fault Sets are defined, you can use any combination of fault units, for example:

- SDS1, SDS2, SDS3
- FS1, SDS1, SDS2
- FS1, FS2, SDS1
- FS1, FS2, FS3
  Figure 3 on page 33 illustrates the same configuration as Figure 2, with the addition of Fault Sets.

Figure 3 Protection Domains, Storage Pools, and Fault Sets

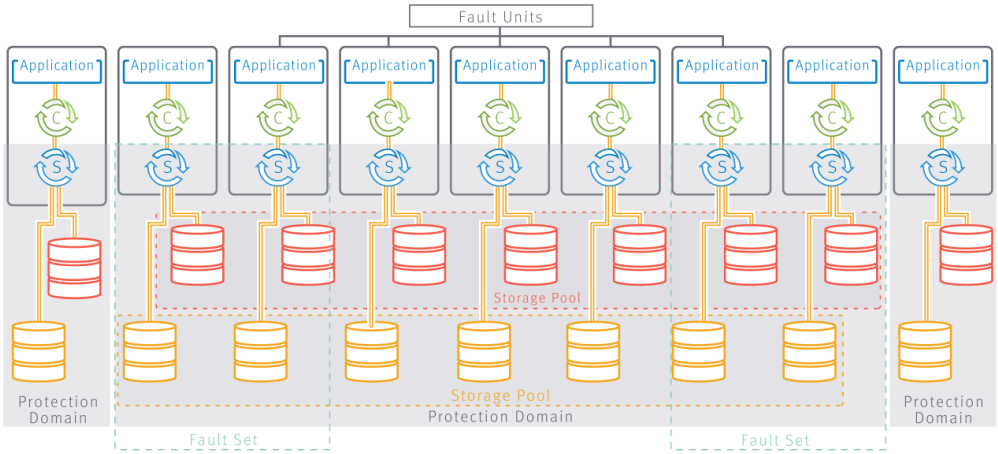

To use Fault Sets, you must work in the following order:

1. Ensure that a Protection Domain exists, or add a new one.
2. Ensure that a Storage Pool and Fault Sets (minimum of three fault units) exist, or add new ones.
3. Add the SDS, designating the Protection Domain and Fault Set, and at the same time, adding the SDS devices into a Storage Pool.
   The automated deployment and installation tools follow this order automatically.

You can only create and configure Fault Sets before adding SDSs to the system, and configuring them incorrectly may prevent the creation of volumes. An SDS can only be added to a Fault Set during the creation of the SDS.

You define Fault Sets and add SDSs to them during installation, using the following management tools:

- Deployment wizard

You can also add Fault Sets when adding SDS nodes after initial installation.

Naming

It is recommended to name all VxFlex OS objects with meaningful names. This will make it easier when defining volumes, associating them with applications, etc.

From the previous example, the Storage Pools can be named "Capacity_Storage" and "Performance_Storage," which allows you to identify the different tiers.

As for Protection Domains, one example would be separating the SDSs used by the finance department from those used by the engineering department. This segregation of different departments is very beneficial in many aspects (security being one of them). Thus, one might name the domains "Financial-PD" and "Engineering-PD."

The Fault Sets could be called "FS_Rack01" and "FS_Rack02."

Protection and load balancing

VxFlex OS maintains the user data in a RAID-1 mesh mirrored layout. Each piece of data is stored on two different servers. The copies are randomly distributed over the storage devices. Rebuild and rebalance processes are fully automated, but are configurable.

Rebuild

When a failure occurs, such as on a server, device or network failure, VxFlex OS immediately initiates a process of protecting the data. This process is called Rebuild, and comes in two flavors:

- Forward rebuild is the process of creating another copy of the data on a new server. In this process, all the devices in the Storage Pool work together, in a many-to-many fashion, to create new copies of all the failed storage blocks. This method ensures an extremely fast rebuild.
- Backward rebuild is the process of re-synchronization of one of the copies. This is done by passing to the copy only changes made to the data while this copy was inaccessible. This process minimizes the amount of data transferred over the network during recovery.

VxFlex OS automatically selects the type of rebuild to perform. This implies that in some cases, more data will be transferred to minimize the time that the user data is not fully protected.

Rebuild throttling

Rebuild throttling sets the rebuild priority policy for a Storage Pool. The policy determines the priority between the rebuild I/O and the application IO when accessing SDS devices. Please note that application I/Os are continuously served.

Applying rebuild throttling will on one hand increase the time the system is exposed with a single copy of some of data, but on the other hand, will reduce the impact on

the application. One has to make a decision and choose the right balance between the two.

The following possible priority policies may be applied:

- No Limit: No limit on rebuild I/Os.

  Any rebuild I/O is submitted to the device immediately, without further queuing. Please note that rebuild I/Os are relatively large and hence setting this policy will speed up the rebuild, but will have the maximal effect on the application I/O.

- Limit Concurrent I/O: Limit the number of concurrent rebuild I/Os per SDS device (default).

  The rebuild I/Os are limited to a predefined number of concurrent I/Os. Once the limit is reached, the next incoming rebuild I/O waits until the completion of a currently executed rebuild I/O. This will complete the Rebuild quickly for best reliability, however, there is a risk of host application impact.

- Favor Application I/O: Limit rebuild in both bandwidth and concurrent I/Os.

  The rebuild I/Os are limited both in bandwidth and in the amount of concurrent I/Os. As long as the number of concurrent rebuild I/Os, and the bandwidth they consume, do not exceed the predefined limits, rebuild I/Os will be served. Once either threshold is reached, the rebuild I/Os wait until both I/O and bandwidth are below their thresholds. For example, setting the value to "1" will guarantee the device will only have one concurrent rebuild IO at any given moment, which will ensure the application IOs only wait for 1 rebuild IO at worst case.

  This imposes bandwidth on top of the Limit Concurrent I/Os option, which is a prerequisite to using this policy.

- Dynamic Bandwidth Throttling: This policy is similar to Favor Application I/O, but extends the interval in which application I/Os are considered to be flowing by defining a minimal quiet period. This quiet period is defined as a certain interval in which no application I/Os occurred. Note that the limits on the rebuild bandwidth and concurrent I/Os are still imposed.

- Default Values:

  - The default policy for rebuild is: Limit Concurrent I/O
  - Rebuild concurrent I/O Limit: 1 concurrent I/O

  Note

  Rebuild throttling affects the system's performance and should only be used by advanced users.

Rebalance

Rebalance is the process of moving one of the data copies to a different server. It occurs when VxFlex OS detects that the user data is not evenly balanced across the devices in a Storage Pool. This can occur as a result of several conditions such as: SDS addition/removal, device addition/removal, or following a recovery from a failure. VxFlex OS will move copies of the data from the most utilized devices to the least utilized ones.

Both Rebuild and Rebalance compete with the application IO for the system resources. This includes network, CPU and disks. VxFlex OS provides a very rich set of parameters that can control this resource consumption. While the system is factory-tuned for balancing between speedy rebuild/rebalance and minimization of

the effect on the application IO, the user has very fine-grain control over the rebuild and rebalance behavior.

Rebalance throttling

Rebalance throttling sets the rebalance priority policy for a Storage Pool. The policy determines the priority between the rebalance I/O and the application IO when accessing SDS devices. Please note that application I/Os are continuously served. Rebalance, unlike rebuild, does not impact the system's reliability and therefore reducing its impact is not risky.

Note

Rebalance throttling affects the system's performance and should only be used by advanced users.

The following possible priority policies may be applied:

- No Limit: No limit on rebalance I/Os.

  Any rebalance I/O is submitted to the device immediately, without further queuing. Please note that rebalance I/Os are relatively large and hence setting this policy will speed up the rebalance, but will have the maximal effect on the application I/O.
- Limit Concurrent I/O: Limit the number of concurrent rebalance I/Os per SDS device.

  The rebalance I/Os are limited to a predefined number of concurrent I/Os. Once the limit is reached, the next incoming rebalance I/O waits until the completion of a currently executed rebalance I/O. For example, setting the value to "1" will guarantee that the device will only have one rebalance IO at any given moment, which will ensure that the application IOs only wait for 1 rebalance IO in the worst case.
- Favor Application I/O: Limit rebalance in both bandwidth and concurrent I/Os.

  The rebalance I/Os are limited both in bandwidth and in the amount of concurrent I/Os. As long as the number of concurrent rebalance I/Os, and the bandwidth they consume, do not exceed the predefined limits, rebalance I/Os will be served. Once either limiter is reached, the rebalance I/Os wait until such time that the limits are not met again.

  This imposes a bandwidth limit on top of the Limit Concurrent I/Os option.
- Dynamic Bandwidth Throttling: This policy is similar to Favor Application I/O, but extends the interval in which application I/Os are considered to be flowing by defining a minimal quiet period. This quiet period is defined as a certain interval in which no application I/Os occurred. Note that the limits on the rebalance bandwidth and concurrent I/Os are still imposed.
- Default Values:
  - The default policy for rebalance: Favor Application I/O
  - Rebalance concurrent I/O Limit: 1 concurrent I/O per SDS device
  - Rebalance bandwidth limit: 10240 KB/s

Checksum protection

This feature addresses errors that change the payload during the transit through VxFlex OS. VxFlex OS protects data in-flight by calculating and validating the checksum value for the payload at both ends.

Note

The checksum feature may have a major impact on performance and availability. Contact Customer Support to verify if your use case is relevant.

- During write operations, the checksum is calculated when the SDC receives the write request from the application. This checksum is validated just before each SDS writes the data on the storage device.
- During read operations, the checksum is calculated when the data is read from the SDS device, and is validated by the SDC before the data returns to the application. If the validating end detects a discrepancy, it will initiate a retry. The checksum will be done in the granularity of a sector (1/2KB).

This feature applies to all IOs: Application, Rebuild, Rebalance, and Migrate. The checksum is also kept in RMcache (Read Memory Cache), protecting every block that is maintained in SDS memory against memory corruption. The checksum feature can be enabled at the Protection Domain level, and defined at the Storage Pool level. The feature is T10/DIF-ready.

ESX vStorage APIs for Array Integration (VAAI)

ESX vStorage APIs for Array Integration (VAAI) is a feature introduced in ESXi/ESX 4.1 that provides hardware acceleration functionality. It allows the host to offload specific virtual machine and storage management operations to compliant storage hardware. With the storage hardware's assistance, the host performs these operations faster, and consumes less CPU, memory, and storage fabric bandwidth.

VAAI uses these fundamental operations:

- Atomic Test & Set (ATS), which is used during creation and locking of files on the VMFS volume
- Clone Blocks/Full Copy/XCOPY, which is used to copy or migrate data within the same physical array
- Zero Blocks/Write Same, which is used to zero-out disk regions
- Thin Provisioning in ESXi 5.x and later hosts, which allows the ESXi host to tell the array when the space previously occupied by a virtual machine (whether it is deleted or migrated to another datastore) can be reclaimed on thin provisioned LUNs.
- Block Delete in ESXi 5.x and later hosts, which allows for space to be reclaimed using the SCSI UNMAP feature.

The VxFlex OS supported VAAI features are:

- Atomic Test & Set (ATS)
- Zero Blocks/Write Same
- Thin Provisioning in ESXi 5.x and later hosts
- Block Delete in ESXi 5.x and later hosts

The following output is an example of typical output:

```
esxcli storage core device vaai status get -d
eui.7dbf14034834bbe01bf7e55800000002
eui.7dbf14034834bbe01bf7e55800000002
VAAI Plugin Name:
ATS Status: supported
```

`Clone Status: unsupported` This means that Clone Block/Full Copy/Xcopy is not supported.

`Zero Status: supported` This means that write same is supported.

`Delete Status: supported` This means that UNMAP is supported.

Note

Thin provisioning is not shown in VAAI output.

Caching

VxFlex OS offers a number of caching options, for the purpose of enhancing system performance.

The following caching options are supported by VxFlex OS:

- RAM Read Cache (using DRAM server memory)
- Read Flash Cache (using SSD and NVMe SSD devices)

In addition, the following caching solution is available, in VxFlex Ready Node systems only:

- DAS Cache (using SSD devices)

Note

DAS Cache is not supported on RHEL 7.4 operating systems nor on PowerEdge 14G servers.

SSDs used for caching cannot be used for storage purposes.

Read Flash Cache (RFcache) and DAS Cache cannot be used simultaneously in the same Storage Pool. However, a node can contain devices using different caching types, as long as the corresponding devices are allocated to different dedicated Storage Pools (one Storage Pool using DAS Cache, and another using Read Flash Cache.

The following table summarizes information about the caching modes provided by the system.

Table 4 Caching modes

Mode	Description	Considerations	Default Setting
RAM Read Cache (RMcache)	Read-only caching performed by server RAM.	RAM Read cache, the fastest type of caching, uses RAM that is allocated for caching. Its size is limited to the amount of allocated RAM. **Note** The amount that may be allocated is limited, and can never be the maximum available RAM.	Disabled, except when storage-only nodes are deployed.
Read Flash Cache (RFcache)	Read-only caching performed by one or more dedicated SSD or NVMe SSD devices in the server.	RFcache uses the full capacity of SSD devices (up to eight) to provide a larger footprint of read-only LRU (Least Recently Used) based-caching resources for the SDS. This type of	

Table 4 Caching modes (continued)

Mode	Description	Considerations	Default Setting
		caching reacts quickly to workload changes to speed up HDD Read performance. Several SSD devices can be allocated to a shared cache pool, and therefore the cache size is limited in size only by the amount of SSDs allocated for this purpose. The RFcache driver must be installed during deployment. Caching devices can be defined either during the installation process or after deployment. Limitations: RFcache does not support partitions on devices installed on Windows nodes. Support matrix: • An RFcache device (flash device) can be partitioned only on Linux. • An SDS storage/source device cannot be partitioned if it needs to be accelerated by RFcache. • An SDS storage/source device as a file (over file system), cannot be accelerated by RFcache.	
DAS Cache	Read and write-back caching performed by one or more dedicated SSD devices in the server	DAS Cache uses the full capacity of one or more SSD devices to provide a large footprint of both read and write-back caching resources to the SDS. This caching mode moves "hot" (active) chunks of data from HDDs to cache, for Read and Write buffering. For write-back caching, the write is temporarily written to the SSD, which is much faster than an HDD, allowing faster response of the SDS to write acknowledgment. One SSD device can accelerate several HDDs (in DAS Cache they are called "Volumes"). Striping the Cache on two devices is not supported in the VxFlex Ready Node solution. **Note** If a fault occurs in the caching device before the writes have been offloaded, all the HDD devices cached by DAS Cache acquire failed status, and a rebuild process commences in VxFlex OS. Once the rebuild is over, the caching disk can be replaced, all caching has stopped in the storage pool, and the HDD members in the storage pool can be cleared of errors.	Disabled

The following table illustrates the caching support matrix:

Table 5 Caching support matrix

System	RFcache	RMcache	DAS Cache	CacheCade
VxFlex OS	Yes	Yes		
VxFlex Ready Node PowerEdge 13G servers	Yes	Yes	Yes	
VxFlex Ready Node PowerEdge 14G servers	Yes	Yes		
VxRack Node 100 Series	Yes	Yes		Yes

Networking

In VxFlex OS, inter-node communication (for the purposes of managing data locations, rebuild and rebalance, and for application access to stored data) can be done on one IP network, or on separate IP networks. VLANS are also supported. Management (via any of the management interfaces) can be done in the following ways:

- Via a separate network with access to the other VxFlex OS components
- On the same network

These options can be configured during deployment, using the deployment wizard.

IP Addresses

The AMS or discovery server must reside on the same Layer 2 physical network (Broadcast Domain). The Ethernet broadcast protocol is used during the initial detection and deployment process.

Prepare IP addresses in your network for the VxFlex OS nodes, based on the following calculations:

Table 6 VxFlex OS management IP network

Item	Description	Comments
N	Number of nodes	
IP address pools	The pools of IP addresses used for static allocation for the following groups: 1. Node_MGMT_IP =VxFlex Ready Nodes' management IP addresses 2. SVM_MGMT_IP =VxFlex OS storage VM management IP address 3. BMC_MGMT_IP =BMC (iDRAC) IP address for each node	SVM IP addresses are relevant for ESXi nodes, only.
AMS server \Proxy IP address	The IP address used by the AMS server \proxy SW to manage the VxFlex OS Nodes and logical entities	Recommended to reside on the same subnet, unless the proxy is used for Discovery operations

Table 6 VxFlex OS management IP network (continued)

Item	Description	Comments
The formula to calculate IP address subnet pool or subnet needs is: N*Node_MGMT_IP +N*SVM_MGMT_IP+N*BMC_MGMT_IP + Optional (AMS server IP address)		

The VxFlex OS data IP network subnet requirements:

- ESXi - 2 separate subnets
- RHEL - 2 or 4 separate subnets

Note

Each pool can belong to a single subnet or to several subnets, as long as the AMS server has routable networking access to the management IP address.

Note

If the AMS is not on the same subnet as the SVM management IP addresses, the discovery server used as a proxy must reside on the same switch network broadcast domain, in order to allow seamless detection and management IP address configuration of the SVM and ESXi IP addresses of all nodes in the system.

VMware limitations

There are limitations in the VMware environment.

Multiple IP subnets
Multiple IP subnets used for the VxFlex OS data network cannot be on the same subnet in a VMware setup.

For more information, see the VMware limitation in the following link:

https://kb.vmware.com/s/article/2010877

Table 7 VxFlex OS data IP network (continued)

Item	Description	Comments
N	Number of nodes	
IP pool for Subnet #1	The pools of IP addresses used for static allocation for the following groups: 1. Node_DATA1_IP = VxFlex Ready Nodes' internal (interconnect) IP addresses 2. SVM_DATA1_IP = VxFlex OS Storage VM Management IP addresses 3. MDM_Cluster_Virtual_IP_DATA1 = the Virtual IP addresses of the MDM cluster in the "Data1" network	For clarity, the first subnet is referred to as "Data1" SVM IP addresses are relevant for ESXi nodes, only.

Table 7 VxFlex OS data IP network (continued) (continued)

Item	Description	Comments
	The formula to calculate the subnet pool or subnet needs is: N*Node_DATA1_IP+N*SVM_DATA1_IP +MDM_Cluster_Virtual_IP_DATA1	
IP pool for Subnet #2	The pools of IP addresses used for static allocation for the following groups: 1. Node_DATA2_IP = VxFlex Ready Nodes' internal (interconnect) IP addresses 2. SVM_DATA2_IP = VxFlex OS Storage VM Management IP addresses 3. MDM_Cluster_Virtual_IP_DATA2 = the Virtual IP addresses of the MDM cluster in the "Data2" network	For clarity, the second subnet is referred to as "Data2" SVM IP addresses are relevant for ESXi nodes, only.
	The formula to calculate the subnet pool or subnet needs is: N*Node_DATA2_IP+N*SVM_DATA2_IP +MDM_Cluster_Virtual_IP_DATA2	

IPv6 addresses

IPv6 is not supported in VMware converged solutions. For more information, contact Customer Support.

VLAN support

VLAN support provides much greater flexibility in system deployment. VLAN IDs can be assigned to various system components (SVM, ESXi Host, BMC (iDRAC), data network 1 and data network 2). Each data and management network can reside on a different VLAN, enabling network separation at a much higher granularity than previously possible.

The network administrator may configure a different VLAN ID for each of the following networks:

- Management ESXi—1 or 10/25 GbE network
- Management SVM—1 or 10/25 GbE network
- Management BMC (iDRAC)—1 or 10 GbE network
- Data A—10/25/100 GbE network
- Data B—10/25/100 GbE network

During deployment, the ESXi image is configured to accept all VLANs by using the VLAN ID 4095 in the ESXi port group. During the Set IP addresses phase, the requested VLAN ID is set in the relevant Port group or vmkernel on the ESXi or the BMC (iDRAC) port using IPMI related commands.

Note

25 and 100 GbE speeds are only supported on 14G platforms.

Networking and connectivity - 13G servers

This section explains network connectivity, cabling best practices, and the most common port configurations for 13G VxFlex Ready Node servers.

Networking connectivity architecture and cabling best practice

This topic describes connectivity architecture, cabling best practice information, and cable connection examples from typical VxFlex Ready Node configurations in order to help you plan your network. Where possible, it is recommended that you use dedicated network adapters for VxFlex OS traffic. However, shared network adapters are fully supported, and might be mandatory in hyper-converged environments.

Note

If you are not familiar with VxFlex Ready Node system architecture, refer to the "System" section in this document.

On ESXi servers, VLANs can be used as long as the broadcast domain is the same between the AMS/Discovery Server and the ESXi management port. For all OS types, the best way to use VLANs with VxFlex Ready Node is to configure the switch port to Mode Access, and assign it to one VLAN. A management switch that has only one use is recommended. Usually, this implies one IP subnet, but the switch can contain several subnets on a single VLAN.

When using several VLANs on the same port, the port must to be configured as Trunk, and you must assign a VLAN ID on each vSwitch on the ESXi. This configuration can be performed on a VxFlex Ready Node server after deployment is complete by updating the VM network and relevant VMKernel port to use the VLAN ID set by the switch manager. This is similar to any other VLAN implementation, and it must be performed after deployment because the AMS deployment flow creates and configures the vSwitches and vmnics as part of the deployment flow.

Note

AMS uses APIPA IP addresses (169.254.x.x/16) for Broadcast and Unicast, and this requires special configuration during deployment.

It is mandatory to set the Management VLAN as either a Layer 2 VLAN without an IP address, or as VLAN 0 (meaning, no VLAN). This will facilitate a seamless onsite "Add Node" operation.

The following figures show how to connect cables to single-node servers in a system consisting of only single-node servers. Data is connected to two switches, for high availability, via the nodes' 10 GbE ports. Management interfaces are connected to a switch on a separate management network using 1 GbE and BMC (iDRAC) ports.

The AMS/Discovery server must be on the same Layer 2 network as the management ports. This allows the use of the broadcast protocol for initial IP address configuration.

The BMC (iDRAC) IP address can reside on a different subnet and/or VLAN, and the only requirement from the BMC (iDRAC) side is that the AMS must have a routable IP Path to the BMC (iDRAC) IP address on the node.

Figure 4 Single-node R730xd (2U1N) server (storage only) typical connection diagram

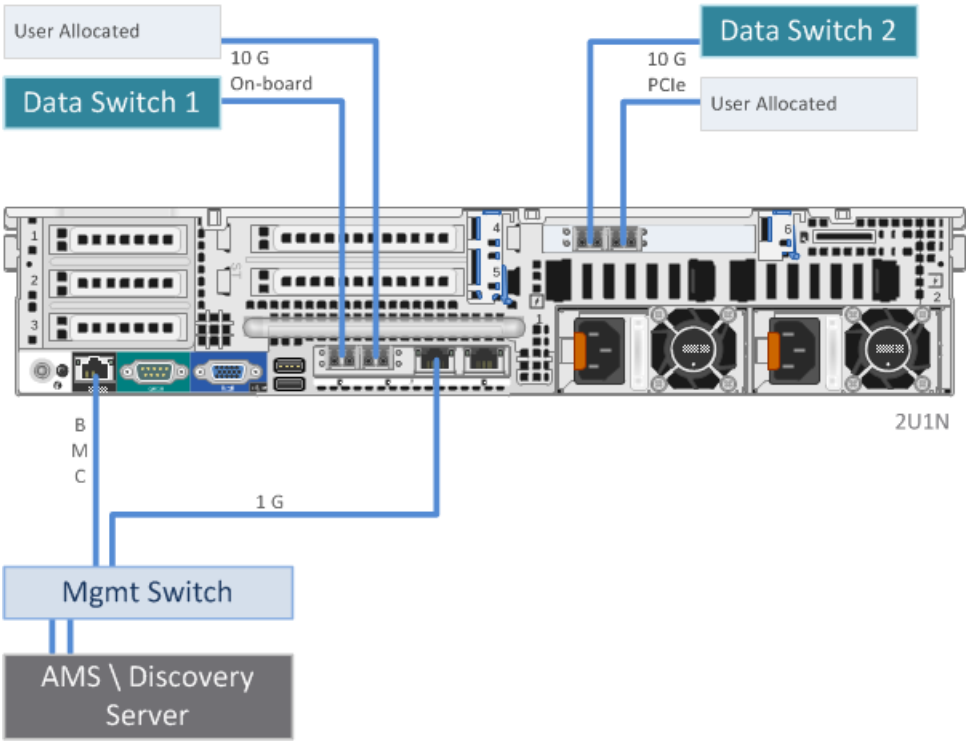

Figure 5 Single-node R730xd (2U1N) server (converged) typical connection diagram

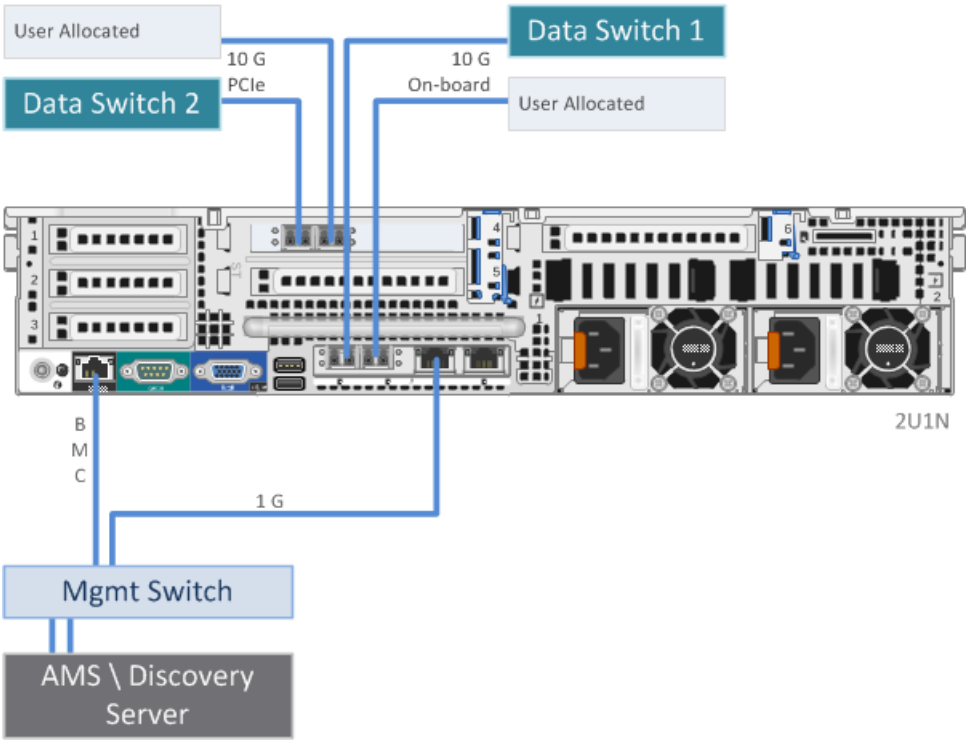

Figure 6 Single-node R630 (1U1N) server typical connection diagram

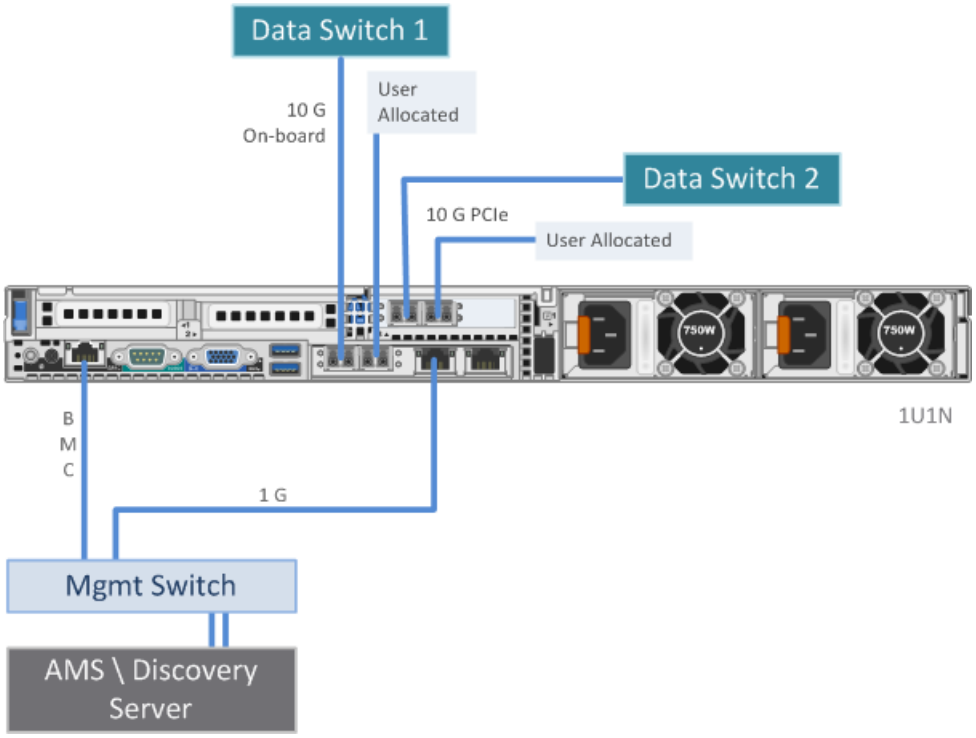

Server node port designations - ESXi 13G servers

For single-node servers in an ESXi-based system, connect the cables and configure the system as shown in the following tables, depending on the server type and architecture, when relevant.

The tables below show the most common correlation between the physical NIC port location and the logical configuration in VMware. When using the same slot locations as shown below, all settings but one setting are static and can be set by the AMS. The vmnic number is the dynamic setting and is determined by the VMware bus scan and boot logic. This number cannot be set persistently.

If you need to build your own table of reference, correlate the ESXi physical ports to the vmnic logical port name.

Note

In this AMS release, Advanced Deployment mode enables you to dynamically select the port to be used in the server. It is recommended, but not mandatory, that you use the VxFlex OS best practice ports. If the VLAN ID is defined, it must also be configured on the switch side connected to the relevant selected port.

Table 8 R730xd single-node server configuration

Description	BMC (iDRAC) Monitoring	VxFlex OS Data1	VxFlex OS Client Network Port1	VxFlex OS Client Network Port2	VxFlex OS Data2	VxFlex OS Managemen t Network	Not in Use
Port	BMC (iDRAC)	vmnic0	vmnic1	vmnic4	vmnic5	vmnic2	vmnic3

Table 8 R730xd single-node server configuration (continued)

Description	BMC (iDRAC) Monitoring	VxFlex OS Data1	VxFlex OS Client Network Port1	VxFlex OS Client Network Port2	VxFlex OS Data2	VxFlex OS Management Network	Not in Use
Physical Port on node	BMC (iDRAC) (remote KVM)	Onboard 10G (Left)	Onboard 10G (Right)	PCI 10G (Right)	PCI 10G (Left)	Onboard (Left)	Onboard 1G (Right)
Speed	1 GbE	10 GbE	10 GbE	10 GbE	10 GbE	1 GbE	1 GbE
vSwitch	N/A	sio_vSwitch1	User allocated	User allocated	sio_vSwitch2	vSwitch0	N/A
VMkernel port		vmk2	User configurable	User configurable	vmk3	vmk0	User configurable
Physical Switch	Mgmt Switch	Switch 1	Switch 1	Switch 2	Switch 2	Mgmt Switch	N/A

Note

The default supported slots for R730xd are slots 4 and 6, as shown in figures Figure 4 on page 44 and Figure 5 on page 44.

Table 9 R630 single-node server configuration

Description	BMC (iDRAC) Monitoring	VxFlex OS Data1	VxFlex OS Client Network Port1	VxFlex OS Client Network Port2	VxFlex OS Data2	VxFlex OS Management Network	Not in Use
Port	BMC (iDRAC)	vmnic0	vmnic1	vmnic4	vmnic5	vmnic2	vmnic3
Physical Port on node	BMC (iDRAC) (remote KVM)	Onboard 10G (Left)	Onboard 10G (Right)	PCI 10G (Right)	PCI 10G (Left)	Onboard 1G (Left)	Onboard 1G (Right)
Speed	1 GbE	10 GbE	10 GbE	10 GbE	10 GbE	1 GbE	1 GbE
vSwitch	N/A	sio_vSwitch1	User allocated	User allocated	sio_vSwitch2	vSwitch0	N/A
VMkernel port		vmk2	User configurable	User configurable	vmk3	vmk0	User configurable
Physical Switch	Mgmt Switch	Switch 1	Switch 1	Switch 2	Switch 2	Mgmt Switch	N/A

Note

The expected slot location for the PCIe card is slot 3.

Server node port designations - Linux 13G servers

For single-node VxFlex Ready Node servers running on Linux, connect cables and configure the ports as shown in the configuration tables, according to the server type.

The following additional information should be taken into account:

- The logical bus address is dynamic and will be allocated by the operating system. Use the slot name to correlate to between the bus address and port name of the interface in the operating system.
- In storage-only configurations, the customer ports can be used as additional data ports.

Note

When the AMS configures the data network ports in Linux, the interface names used are sio_d_1 and sio_d_2.

Table 10 R630 single-node server configuration

Description	BMC (iDRAC) Monitoring	VxFlex OS Data1	VxFlex OS Client Network Port1	VxFlex OS Client Network Port2	VxFlex OS Data2	VxFlex OS Management Network	Not in Use
Interface name (RHEL7)	N/A	enoXX	enoXX	enpXsXfX	enpXsXfX	enoXX	enoXX
Slot name	N/A	Integrated NIC1	Integrated NIC2	PCIe Slot3 Port 1	PCIe Slot3 Port 0	Integrated NIC 3	Integrated NIC 4
Logical bus address	N/A	Z:0	Z:1	Y:1	Y:0	X:0	X:1
Physical port on node	BMC (iDRAC)	Onboard 10G (Left)	Onboard 10G (Right)	PCI 10G (Right)	PCI 10G (Left)	Onboard 1G (Left)	Onboard 1G (Right)
Speed	1 GbE	10 GbE	10 GbE	10 GbE	10 GbE	1 GbE	1 GbE
Physical switch	Mgmt Switch	Data Switch1	Data Switch1	Data Switch2	Data Switch2	Mgmt Switch	Not in use

Table 11 R730xd single-node server configuration

Description	BMC (iDRAC) Monitoring	VxFlex OS Data1	VxFlex OS Client Network Port1	VxFlex OS Client Network Port2	VxFlex OS Data2	VxFlex OS Management Network	Not in Use
Interface name (RHEL 7)	N/A	enoXX	enoXX	enpXsXfX	enpXsXfX	enoXX	enoXX
Slot name	N/A	Integrated NIC1	Integrated NIC2	PCIe Slot4 Port 1	PCIe Slot4 Port 0	Integrated NIC 3	Integrated NIC 4
Logical bus address	N/A	Z:0	Z:1	Y:1	Y:0	X:0	X:1
Physical port on node	BMC (iDRAC)	Onboard 10G (Left)	Onboard 10G (Right)	PCI 10G (Right)	PCI 10G (Left)	Onboard 1G (Left)	Onboard 1G (Right)
Speed	1 GbE	10 GbE	10 GbE	10 GbE	10 GbE	1 GbE	1 GbE
Physical Switch	Mgmt Switch	Data Switch1	Data Switch1	Data Switch2	Data Switch2	Mgmt Switch	Not in use

Table 12 R730xd single-node server configuration (Storage Only)

Description	BMC (iDRAC) Monitoring	VxFlex OS Data1	VxFlex OS Client Network Port1	VxFlex OS Client Network Port2	VxFlex OS Data2	VxFlex OS Management Network	Not in Use
Interface name (RHEL 7)	N/A	enoXX	enoXX	enpXsXfX	enpXsXfX	enoXX	enoXX
Slot name	N/A	Integrated NIC1	Integrated NIC2	PCIe Slot6 Port 0	PCIe Slot6 Port 1	Integrated NIC 3	Integrated NIC 4
Logical bus address	N/A	Z:0	Z:1	X:1	X:0	Y:0	Y:1
Physical port on node	BMC (iDRAC)	Onboard 10G (Left)	Onboard 10G (Right)	PCI 10G (Right)	PCI 10G (Left)	Onboard 1G (Left)	Onboard 1G (Right)
Speed	1 GbE	10 GbE	10 GbE	10 GbE	10 GbE	1 GbE	1 GbE
Physical Switch	Mgmt Switch	Data Switch1	Data Switch1	Data Switch2	Data Switch2	Mgmt Switch	Not in use

Correlating PCI slot locations and interface names

Gather and correlate PCI slot locations and interface names in order to build a table with the relevant interface names for Linux nodes.

RHEL 7 does not use the `70-persistent-net.rules` file for Ethernet persistency. Instead, in RHEL 7 nodes, the interface names are in the following format: enXXX (for example: enp6s0f0 or eno3).

Note

In the following procedure, the commands are being run on a RHEL 7 node, so the example outputs include interface names in the format enXXX.

Procedure

1. Find the correlation between physical slot names and interface names so that you can discover the correlation between PCI slot locations and interface names in the operating system:

   Note

   For each of the following commands, write down the output and set it aside for use later in this task:

   a. Correlate the physical PCIe slot to a logical bus address:

   ```
   dmidecode -t slot |egrep "Designation|Current Usage|Bus
   Address"
   ```

   This command displays the PCI location of each PCIe card. In the following sample output, the logical bus address in the last three lines is "0000:05:00.0," which is shown as `Current Usage: In Use` and correlates to physical PCIe card slot location "PCIe Slot 6."

In this case, the 10GB PCI NIC is presented on slot 6, and the logical bus address of its ports is 5:0 and 5:1. (See Figure 7 on page 50, in which the correlation between PCIe slot and the logical bus address is highlighted in yellow.)

```
        Designation: PCIe Slot 1
        Current Usage: Available
        Designation: PCIe Slot 2
        Current Usage: Available
        Designation: PCIe Slot 3
        Current Usage: Available
        Designation: PCIe Slot 4
        Current Usage: Available
        Designation: PCIe Slot 5
        Current Usage: Available
        Designation: PCIe Slot 6
        Current Usage: In Use
        Bus Address: 0000:05:00.0
```

b. Correlate the onboard device name (NIC1 to NIC4) to a logical bus address:

```
dmidecode -t baseboard | sed -n -e '/NIC/,/Bus Address/p' |
egrep "NIC|Bus Address"
```

This command displays the onboard location of each NIC port, In the following sample output, the logical bus address in the first two lines is "0000:01:00.0" which correlates to physical NIC `Integrated NIC 1`.

```
        Reference Designation: Integrated NIC 1
        Bus Address: 0000:01:00.0
        Reference Designation: Integrated NIC 2
        Bus Address: 0000:01:00.1
        Reference Designation: Integrated NIC 3
        Bus Address: 0000:09:00.0
        Reference Designation: Integrated NIC 4
        Bus Address: 0000:09:00.1
```

c. Correlate the list of OS interface names to a logical bus address:

```
ip a | grep ": " | awk '{print $2}' | tr -d ":" | grep -v
^"lo"$ | xargs -I '{}' sh -c 'echo {}; ethtool -i {} | grep
bus-info'
```

This command displays the Ethernet network names (enXXX) of the interfaces and the correlating logical bus address (without the loopback). For example:

```
        eno3
        bus-info: 0000:09:00.0
        eno4
        bus-info: 0000:09:00.1
        enp6s0f0
        bus-info: 0000:05:00.0
        enp6s0f1
        bus-info: 0000:05:00.1
        eno1
        bus-info: 0000:01:00.0
        eno2
        bus-info: 0000:01:00.1
```

Figure 7 Correlating physical slot and Ethernet network names

A

```
Designation: PCIe Slot 1
 Current Usage: Available
 Designation: PCIe Slot 2
 Current Usage: Available
 Designation: PCIe Slot 3
 Current Usage: Available
 Designation: PCIe Slot 4
 Current Usage: Available
 Designation: PCIe Slot 5
 Current Usage: Available
 Designation: PCIe Slot 6
 Current Usage: In Use
 Bus Address: 0000:05:00.0
```

B

```
Reference Designation: Integrated NIC 1
 Bus Address: 0000:01:00.0
 Reference Designation: Integrated NIC 2
 Bus Address: 0000:01:00.1
 Reference Designation: Integrated NIC 3
 Bus Address: 0000:09:00.0
 Reference Designation: Integrated NIC 4
 Bus Address: 0000:09:00.1
```

C

```
eno3
bus-info: 0000:09:00.0
eno4
bus-info: 0000:09:00.1
enp6s0f0
bus-info: 0000:05:00.0
enp6s0f1
bus-info: 0000:05:00.1
eno1
bus-info: 0000:01:00.0
eno2
bus-info: 0000:01:00.1
```

2. Create a table for your own use, similar to the example table shown below:
   - Base the table on the RHEL port definitions for the relevant server type described in Server node port designations - Linux 13G servers on page 46.
   - In the **Port** row, record the enXXX names that correspond with the PCI slots.
   - Map each line in turn from Step 1.a on page 48 output to the corresponding line in Step 1.b on page 49 output.

Note

In 1U1N and 2U1N platforms, the 1 GB port on the right is used for management, and the left 10 GB ports on both NICs are used for data.

For example, the last column in the following table was created using the information from the 4th line in steps 1.a on page 48 and 1.b on page 49:

VxFlex OS role	...	Optional VxFlex OS Data2 2nd port
Port		enp6s0f0
Slot name		PCIe Slot 6 Port 0
Logical bus address		0000:05:00.0
Physical port on node		PCI 10G Left
Speed		0 GB

Networking and connectivity - 14G servers

This section explains network connectivity, cabling best practices, and the most common port configurations for 14G VxFlex Ready Node servers.

Networking connectivity architecture and cabling best practice (R640, R740XD servers)

This topic describes connectivity architecture, cabling best practice information, and cable connection examples from typical VxFlex Ready Node configurations in order to help you plan your network.

Note

If you are not familiar with VxFlex Ready Node system architecture, refer to the "System" section in this document.

On ESXi servers, VLANs can be used as long as the broadcast domain is the same between the AMS/Discovery Server and the ESXi management port. For all OS types,

the best way to use VLANs with VxFlex Ready Node is to configure the switch port to Mode Access, and assign it to one VLAN. A management switch that has only one use is recommended. Usually, this implies one IP subnet, but the switch can contain several subnets on a single VLAN.

When using several VLANs on the same port, the port must to be configured as Trunk, and you must assign a VLAN ID on each vSwitch on the ESXi. This configuration can be performed on a VxFlex Ready Node server after deployment is complete by updating the VM network and relevant VMKernel port to use the VLAN ID set by the switch manager. This is similar to any other VLAN implementation, and it must be performed after deployment because the AMS deployment flow creates and configures the vSwitches and vmnics as part of the deployment flow.

Note

AMS uses APIPA IP addresses (169.254.x.x/16) for broadcast and unicast, and this requires special configuration during deployment.

It is mandatory to set the management VLAN as either a Layer 2 VLAN without an IP address, or as VLAN 0 (meaning, no VLAN). This facilitates a seamless onsite "Add Node" operation.

R640 and R740xd server NIC and BMC (iDRAC) port locations

NIC and BMC (iDRAC) port designations for VxFlex Ready Node R640 and R740xd servers for various server configurations are described in this topic. NICs may be located onboard, or in PCI slots, depending on the server configuration.

The following figure shows possible connector locations on the rear panel of the servers. Actual NIC locations depend on server configurations. Some models do not contain NICs in all possible slots. Items marked 1 to 8 represent PCI slots.

Figure 8 R640 and R740xd PCI slots, integrated NICs and BMC port locations

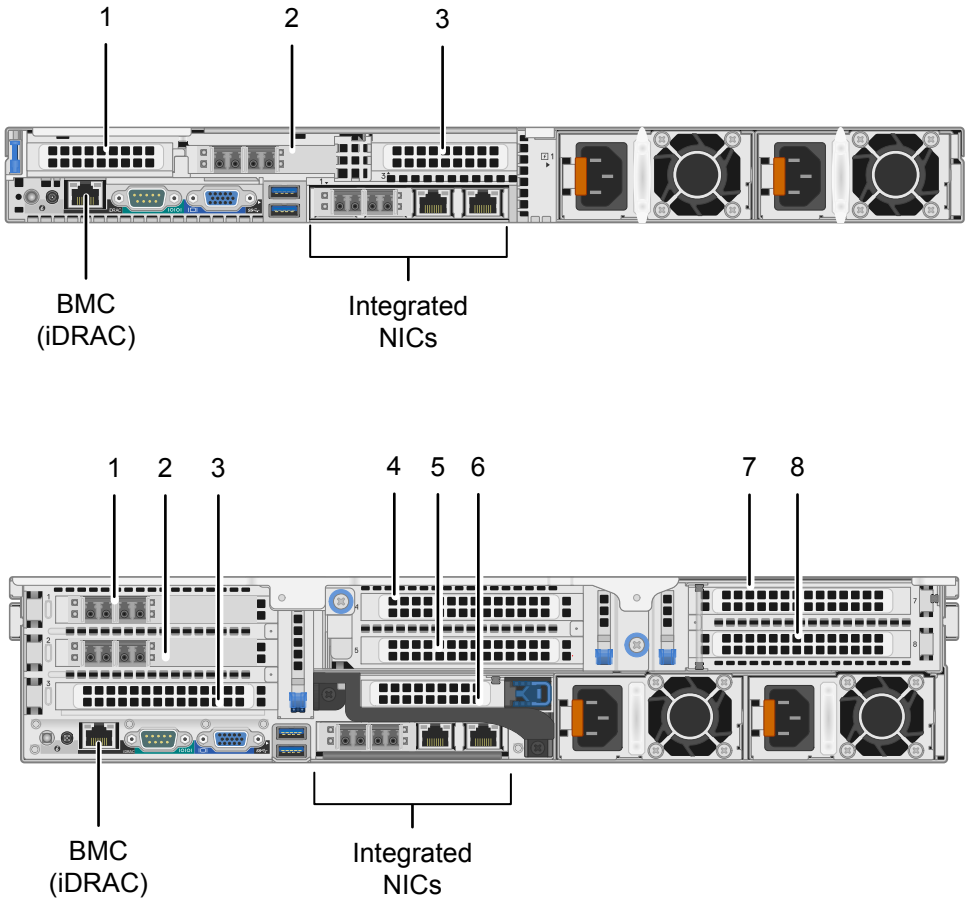

Data cables are connected to two switches for high availability, via the nodes' 10 GbE, 25 GbE, or 100 GbE ports. Management interfaces are connected to a switch on a separate management network using onboard and BMC (iDRAC) ports.

In ESXi environments, the AMS/discovery server must be on the same Layer 2 network as the ESXi management ports. This allows the use of the Broadcast protocol for initial IP address configuration.

The BMC (iDRAC) IP address can reside on a different subnet and/or VLAN, and the only requirement from the BMC (iDRAC) side is that the AMS (when used) must have a routable IP Path to the BMC (iDRAC) IP address on the node.

The default configuration used in simple mode deployment follows the following rules of thumb:

- If there is a 1G NIC onboard, the left one of the pair is always used for the management network.
- If there are two PCI cards or more, the lowest slot number is used for allocation of data ports. For example, if slots 2, 3 and 5 are populated with NICs, AMS uses slots 2 and 3, and the left port on each PCI card is used for data.

  Note

  You can always change the default using advanced deployment mode in AMS.

- If there is only one NIC, the left 10G onboard port is used for the data network.

- If there are only 2*25GbE onboard ports, the left port is used for both the management and data network, and the right port for is used for data.

  Note

  The application\client network will have to run on those same two ports.

- If there are only 4*10G onboard ports:
  - The left port is used for management
  - The left port, and the third port from the left are used for VxFlex OS data
  - The second port from the left, and the right port are used for application\client traffic

VxFlex Ready Node R640 and R740xd port designations - ESXi

Connect the cables and configure the system in an ESXi-based system as shown in the following tables, depending on the server type and architecture, when relevant.

The tables below show the most common correlation between the physical NIC port location and the logical configuration in VMware. When using the same slot locations as shown below, all settings but one setting are static and can be set by the AMS. The vmnic number is the dynamic setting and is determined by the VMware bus scan and boot logic. This number cannot be set persistently.

If you need to build your own table of reference, correlate the ESXi physical ports to the vmnic logical port name.

Note

In this AMS release, Advanced Deployment mode enables you to dynamically select the port to be used in the server. It is recommended, but not mandatory, that you use the VxFlex OS best practice ports. If the VLAN ID is defined, it must also be configured on the switch side connected to the relevant selected port.

R740xd servers

Note

In the following tables use SFP+ for 10 GbE, and SFP28 for 25 GbE ports.

Table 13 R740xd server configuration for 10/25 GbE, 32xSAS, SFP+/SFP28

Description	BMC (iDRAC) Monitoring	VxFlex OS Data1	VxFlex OS Client Network Port1	VxFlex OS Client Network Port2	VxFlex OS Data2	VxFlex OS Managemen t Network	Not in Use
Port	BMC (iDRAC)	vmnic0	vmnic1	vmnic4	vmnic5	vmnic2	vmnic3
PCI slot	N/A	1:0	1:1	2:1	2:0	N/A	N/A
Physical Port on node	BMC (iDRAC) (remote KVM)	PCI 10/25 GbE (Left)	PCI 10/25 GbE (Right)	PCI 10/25 GbE (Right)	PCI 10/25 GbE (Left)	Onboard 1 GbE (Left)	N/A
Speed	1 GbE	10/25 GbE	10/25 GbE	10/25 GbE	10/25 GbE	1 GbE	N/A

Table 13 R740xd server configuration for 10/25 GbE, 32xSAS, SFP+/SFP28 (continued)

Description	BMC (iDRAC) Monitoring	VxFlex OS Data1	VxFlex OS Client Network Port1	VxFlex OS Client Network Port2	VxFlex OS Data2	VxFlex OS Managemen t Network	Not in Use
vSwitch	N/A	sio_vSwitch1	User allocated	User allocated	sio_vSwi tch2	vSwitch0	N/A
VMkernel port		vmk2	User configurable	User configurable	vmk3	vmk0	N/A
Physical Switch	Mgmt Switch	Switch 1	Switch 1	Switch 2	Switch 2	Mgmt Switch	N/A

Note

The above configuration supports more than one PCIe slot-based NIC.

Table 14 R740xd server configuration for 10/25GbE, 12xNVMe, SFP+/SFP28

Description	BMC (iDRAC) Monitoring	VxFlex OS Data1	VxFlex OS Client Network Port1	VxFlex OS Client Network Port2	VxFlex OS Data2	VxFlex OS Managemen t Network	Not in Use
Port	BMC (iDRAC)	vmnic0	vmnic1	vmnic4	vmnic5	vmnic2	vmnic3
PCI slot	N/A	2:0	2:1	7:1	7:0	N/A	N/A
Physical Port on node	BMC (iDRAC) (remote KVM)	PCI 10/25 GbE (Left)	PCI 10/25 GbE (Right)	PCI 10/25 GbE(Right)	PCI 10/25 GbE (Left)	Onboard 1 GbE (Left)	N/A
Speed	1 GbE	10/25 GbE	10/25 GbE	10/25 GbE	10/25 GbE	1 GbE	N/A
vSwitch	N/A	sio_vSwitch1	User allocated	User allocated	sio_vSwi tch2	vSwitch0	N/A
VMkernel port		vmk2	User configurable	User configurable	vmk3	vmk0	N/A
Physical Switch	Mgmt Switch	Switch 1	Switch 1	Switch 2	Switch 2	Mgmt Switch	N/A

Note

The above configuration supports more than one PCIe slot-based NIC.

Table 15 R740xd server configuration for 25GbE, 3xGPU (M60)/ 2xGPU (M10), SFP28

Description	BMC (iDRAC) Monitoring	VxFlex OS Data1/ VxFlex OS Management Network	VxFlex OS Data2	Not in Use
Port	BMC (iDRAC)	vmnic1	vmnic0	vmnic3
PCI slot	N/A	N/A	N/A	N/A
Physical Port on node	BMC (iDRAC) (remote KVM)	Onboard 25 GbE (Left)	Onboard 25 GbE (Right)	None
Speed	1 GbE	25 GbE	25 GbE	None
vSwitch	N/A	vSwitch0	sio_vSwitch1	N/A
VMkernel port		vmk2	vmk3	N/A
Physical Switch	Mgmt Switch	Data Switch 1	Data Switch 2	N/A

Note

The above configuration supports only two 25 GbE onboard NICs in three-GPU configuration. In configurations with only two GPUs, another PCIe-based NIC is supported: one in slot 3 and/or one in slot 4.

Table 16 R740xd server configuration for 10GbE, 3xGPU (M60)/ 2xGPU (M10), SFP+

Description	BMC (iDRAC) Monitoring	VxFlex OS Data1	VxFlex OS Client Network Port1	VxFlex OS Client Network Port2	VxFlex OS Data2	VxFlex OS Managemen t Network	Not in Use
Port	BMC (iDRAC)	vmnic0	vmnic1	vmnic3	vmnic2	vmnic0	N/A
PCI slot	N/A	N/A	N/A	N/A	N/A	N/A	N/A
Physical Port on node	BMC (iDRAC) (remote KVM)	Onboard 10 GbE (far Left)	Onboard 10 GbE (second from Left)	Onboard 10 GbE (far Right)	Onboard 10 GbE (second from Right)	Onboard 10 GbE (far Left) - The same as for VxFlex OS Data1	None
Speed	1 GbE	10 GbE	10 GbE	10 GbE	10 GbE	10 GbE	None
vSwitch	N/A	vSwitch0	User allocated	User allocated	sio_vSwi tch1	vSwitch0	N/A
VMkernel port		vmk2	User configurable	User configurable	vmk3	vmk0	N/A
Physical Switch	Mgmt Switch	Data Switch 1	Data Switch 1	Data Switch 2	Data Switch 2	Mgmt/Data1 Switch	N/A

Note

In configurations with only two GPUs, another PCIe-based NIC is supported in either slot 3 or slot 4.

R640 servers

Note

In the following tables use SFP+ for 10 GbE, and SFP28 for 25 GbE ports.

Table 17 R640 server configuration for 10/25GbE, 1CPU, SFP+/SFP28

Description	BMC (iDRAC) Monitoring	VxFlex OS Data1	VxFlex OS Client Network Port1	VxFlex OS Client Network Port2	VxFlex OS Data2	VxFlex OS Managemen t Network	Not in Use
Port	BMC (iDRAC)	vmnic0	vmnic1	vmnic4	vmnic5	vmnic2	N/A
PCI slot	N/A	2:0	2:1	N/A	N/A	N/A	N/A
Physical Port on node	BMC (iDRAC) (remote KVM)	PCI 10/25 GbE (Left)	PCI 10/25 GbE (Right)	Onboard 25 GbE (Right)	Onboard 25 GbE (Left)	Onboard 25 GbE (Left) - The same as for VxFlex OS Data2	N/A
Speed	1 GbE	10/25 GbE	10/25 GbE	25 GbE	25 GbE	25 GbE	N/A
vSwitch	N/A	sio_vSwitch1	User allocated	User allocated	vSwitch1	vSwitch0	N/A
VMkernel port		vmk2	User configurable	User configurable	vmk3	vmk0	N/A
Physical Switch	Mgmt Switch	Switch 1	Switch 1	Switch 2	Switch 2	Mgmt Switch/ Switch 2	N/A

Table 18 R640 server configuration for 10/25GbE, 2CPU, SFP+/SFP28

Description	BMC (iDRAC) Monitoring	VxFlex OS Data1	VxFlex OS Client Network Port1	VxFlex OS Client Network Port2	VxFlex OS Data2	VxFlex OS Managemen t Network	Not in Use
Port	BMC (iDRAC)	vmnic0	vmnic1	vmnic4	vmnic5	vmnic2	N/A
PCI slot	N/A	N/A	N/A	N/A	N/A	N/A	N/A
Physical Port on node	BMC (iDRAC) (remote KVM)	PCI 10/25 GbE (Left)	PCI 10/25 GbE (Right)	PCI 10/25 GbE (Right)	PCI 10/25 GbE (Left)	Onboard 1 GbE (Left)	N/A
Speed	1 GbE	10/25 GbE	10/25 GbE	10/25 GbE	10/25 GbE	1 GbE	N/A
vSwitch	N/A	sio_vSwitch1	User allocated	User allocated	sio_vSwi tch2	vSwitch0	N/A
VMkernel port		vmk2	User configurable	User configurable	vmk3	vmk0	N/A
Physical Switch	Mgmt Switch	Switch 1	Switch 1	Switch 2	Switch 2	Mgmt Switch	N/A

Note

This configuration supports more than one PCIe-based NIC in slots 2 and 3.

VxFlex Ready Node R640 and R740xd port designations - Linux

For single-node VxFlex Ready Node servers running Linux, connect the cables, and configure the ports as shown in the configuration tables according to the server type.

The following additional information should be taken into account:

- The logical bus address is dynamic and will be allocated by the operating system. Use the slot name to correlate to between the bus address and port name of the interface in the operating system.
- In storage-only configurations, the customer ports can be used as additional data ports.

Table 19 R640 single-node server configuration

Description	BMC (iDRAC) Monitoring	VxFlex OS Data1	VxFlex OS Client Network Port1	VxFlex OS Client Network Port2	VxFlex OS Data2	VxFlex OS Management Network	Not in Use
Interface name (RHEL7)	N/A	emXX	emXX	pXpX	pXpX	emXX	emXX
Slot name	N/A	Integrated NIC1	Integrated NIC2	PCIe Slot3 Port 1	PCIe Slot3 Port 0	Integrated NIC 3	Integrated NIC 4
Logical bus address	N/A	Z:0	Z:1	Y:1	Y:0	X:0	X:1
Physical port on node	BMC (iDRAC)	Onboard 10 GbE Left	Onboard 10 GbE Right	PCI 10/25 GbE Right	PCI 10/25 GbE Left	Onboard 1 GbE Left	Onboard 1 GbE Right
Speed	1 GbE	10 GbE	10 GbE	10/25 GbE	10/25 GbE	1 GbE	1 GbE
Physical switch	Mgmt Switch	Data Switch1	Data Switch1	Data Switch2	Data Switch2	Mgmt Switch	Not in use

Note

When using the onboard 10 GbE left and right ports (x710 Intel NICs) in the above configuration, contact Customer Support for guidance.

Note

You can also use a second NIC in Slot 2 for a 4*25 GbE option.

Table 20 R740xd single-node server configuration for 10/25GbE, 2CPU, SFP+/SFP28

Description	BMC (iDRAC) Monitoring	VxFlex OS Data1	VxFlex OS Client Network Port1	VxFlex OS Client Network Port2	VxFlex OS Data2	VxFlex OS Management Network	Not in Use
Interface name (RHEL 7)	N/A	emXX	emXX	pXpX	pXpX	emXX	emXX
Slot name	N/A	Free PCI slot, left port	Free PCIe slot, right port	Free PCIe slot, right port	Free PCIe slot, left port	Integrated NIC 3	Integrated NIC 4
Logical bus address	N/A	Z:0	Z:1	Y:1	Y:0	X:0	X:1
Physical port on node	BMC (iDRAC)	PCI 10/25 GbE Left	PCI 10/25 GbE Right	PCI 10/25 GbE Right	PCI 10/25 GbE Left	Onboard 1 GbE Left	Onboard 1G Right
Speed	1GbE	10/25 GbE	10/25 GbE	10/25 GbE	10/25 GbE	1 GbE	1 GbE
Physical Switch	Mgmt Switch	Data Switch1	Data Switch1	Data Switch2	Data Switch2	Mgmt Switch	Not in use

Note

Two or more PCIe slot-based NICs are supported.

Correlate the PCI slot locations and interface names

Gather and correlate PCI slot locations and interface names in order to build a table with the relevant interface names for the Linux nodes.

RHEL 7 does not use the `70-persistent-net.rules` file for Ethernet persistency. Instead, in RHEL 7 nodes, the interface names are in the following format: enXXX (for example: enp6s0f0 or eno3).

Procedure

1. Find the correlation between physical slot names and interface names so that you can discover the correlation between PCI slot locations and interface names in the operating system:

   Note

   For each of the following commands, write down the output and set it aside for use later in this task:

   a. Correlate the physical PCIe slot to a logical bus address:

   ```
   dmidecode -t slot |egrep "Designation|Current Usage|Bus
   Address"
   ```

   This command displays the PCI location of each PCIe card. In the following sample output, the logical bus address in the last three lines is "0000:05:00.0," which is shown as `Current Usage: In Use` and correlates to physical PCIe card slot location "PCIe Slot 6."

In this case, the 10GB PCI NIC is presented on slot 6, and the logical bus address of its ports is 5:0 and 5:1. (See the "Correlating physical slot and Ethernet network names" figure, in which the correlation between PCIe slot and the logical bus address is highlighted in yellow.)

```
        Designation: PCIe Slot 1
        Current Usage: Available
        Designation: PCIe Slot 2
        Current Usage: Available
        Designation: PCIe Slot 3
        Current Usage: Available
        Designation: PCIe Slot 4
        Current Usage: Available
        Designation: PCIe Slot 5
        Current Usage: Available
        Designation: PCIe Slot 6
        Current Usage: In Use
        Bus Address: 0000:05:00.0
```

b. Correlate the onboard device name (NIC1 to NIC4) to a logical bus address:

```
dmidecode -t baseboard | sed -n -e '/NIC/,/Bus Address/p' |
egrep "NIC|Bus Address"
```

This command displays the onboard location of each NIC port, In the following sample output, the logical bus address in the first two lines is "0000:01:00.0" which correlates to physical NIC `Integrated NIC 1`.

```
        Reference Designation: Integrated NIC 1
        Bus Address: 0000:01:00.0
        Reference Designation: Integrated NIC 2
        Bus Address: 0000:01:00.1
        Reference Designation: Integrated NIC 3
        Bus Address: 0000:09:00.0
        Reference Designation: Integrated NIC 4
        Bus Address: 0000:09:00.1
```

c. Correlate the list of OS interface names to a logical bus address:

```
ip a | grep ": " | awk '{print $2}' | tr -d ":" | grep -v
^"lo"$ | xargs -I '{}' sh -c 'echo {}; ethtool -i {} | grep
bus-info'
```

This command displays the Ethernet network names (enXXX) of the interfaces and the correlating logical bus address (without the loopback). For example:

```
        em1
        bus-info: 0000:19:00.0
        p1p1
        bus-info: 0000:3b:00.0
        em2
        bus-info: 0000:19:00.1
        em3
        bus-info: 0000:19:00.2
        em4
        bus-info: 0000:19:00.3
        p1p2
        bus-info: 0000:3b:00.1
```

Figure 9 Correlating physical slot and Ethernet network names

A

```
Designation: PCIe Slot 1
 Current Usage: Available
 Designation: PCIe Slot 2
 Current Usage: Available
 Designation: PCIe Slot 3
 Current Usage: Available
 Designation: PCIe Slot 4
 Current Usage: Available
 Designation: PCIe Slot 5
 Current Usage: Available
 Designation: PCIe Slot 6
 Current Usage: In Use
 Bus Address: 0000:05:00.0
```

B

```
Reference Designation: Integrated NIC 1
 Bus Address: 0000:01:00.0
 Reference Designation: Integrated NIC 2
 Bus Address: 0000:01:00.1
 Reference Designation: Integrated NIC 3
 Bus Address: 0000:09:00.0
 Reference Designation: Integrated NIC 4
 Bus Address: 0000:09:00.1
```

C

```
eno3
bus-info: 0000:09:00.0
eno4
bus-info: 0000:09:00.1
enp6s0f0
bus-info: 0000:05:00.0
enp6s0f1
bus-info: 0000:05:00.1
eno1
bus-info: 0000:01:00.0
eno2
bus-info: 0000:01:00.1
```

2. Create a table for your own use, similar to the example table shown below:
   - Base the table on the RHEL port definitions for the relevant server type described in the "Server node port designations" section.
   - In the **Port** row, record the enXXX names that correspond with the PCI slots.
   - Map each line in turn from step 1.a on page 58 output to the corresponding line in step 1.b on page 59 output.

Note

In R640 (1U1N) and R740xd (2U1N) platforms, the 1 GbE port on the right is used for management.

For example, the last column in the following table was created using the information from the 4th line in steps 1.a on page 58 and 1.b on page 59:

VxFlex OS role	...	Optional VxFlex OS Data2 2nd port
Port		enp6s0f0
Slot name		PCIe Slot 6 Port 0
Logical bus address		0000:05:00.0
Physical port on node		PCI 10G Left
Speed		0 GB

Virtual IP Address

Virtual IP addresses can be defined for the MDM cluster.

Up to two virtual IP addresses can be defined for the MDM cluster in AMS deployments. SDCs are then mapped to the MDM cluster's virtual IP addresses, instead of to static MDM IP addresses. MDMs are sometimes switched during normal operation of the cluster, and the virtual IP address will always be mapped to the active MDM. The use of virtual IP addresses simplifies maintenance procedures on the MDM cluster, because system components communicate via the virtual IP addresses. Therefore, SDCs do not need to be reconfigured when a server hosting an MDM is replaced.

Existing systems may be extended to include additional MDMs to a cluster. The new MDMs should be mapped to the existing virtual IP addresses.

If virtual IP addresses need to be modified, you must use the CLI or the REST API, and it must be done with extreme caution.

All SDCs will require reconfiguration, to reflect the changes made to the MDM cluster. Otherwise, the SDCs will not be able to communicate with the MDM cluster, and volumes will not be accessible.

Monitoring of SDC and SDS connections

The system monitors all connections between SDCs and SDSs and sends out an alert when an active connection between an SDC and an SDS goes down.

To effectively monitor SDC and SDS connections, the MDM collects connectivity updates from all of the SDCs. The MDM posts events whenever an SDC connects to or disconnects from a specific SDS IP address. The MDM frequently analyzes the connectivity status to determine the current system state. The system does not send out alerts for temporary connectivity issues that are resolved in less than 10 seconds.

The following are the possible connectivity states between SDCs and SDSs in the system:

- All connected
- One SDC is disconnected from one SDS
- One SDC is disconnected from one SDS IP address
- One SDC is disconnected from all SDSs
- All SDCs are disconnected from one SDS
- All SDCs are disconnected from one SDS IP address
- All SDCs are disconnected from all SDSs
- Multiple disconnections

When the system's connectivity state changes to any state other than `All Connected`, an alert is displayed in the VxFlex OS GUI and is written to the MDM event log. Once an alert is generated, you can use the SCLI to query details on the disconnection using the command `scli --query_sdc_to_sds_disconnections`. For more information about running SCLI commands, see the *CLI Reference Guide*.

The MDM does not monitor the connectivity state of SDCs or SDSs in the following scenarios:

- SDS is in maintenance mode
- SDS is disconnected from the MDM
- SDS is in the process of being removed
- SDC is disconnected from the MDM for more than two minutes
- SDC is not approved

The alerts are detailed in "Alerts in SNMP, VxFlex OS GUI, REST, and ESRS" in the *VxFlex OS User Guide*.

Snapshots

The VxFlex OS storage system enables you to take snapshots of existing volumes, up to 31 per volume. The snapshots are thin provisioned and are extremely quick. For more information about thin provisioning, see SAN virtualization layer.

Once a snapshot is generated, it becomes a new, unmapped volume in the system. You can manipulate it in the same manner as any other volume exposed by the VxFlex OS storage system.

Figure 10 Snapshot operations

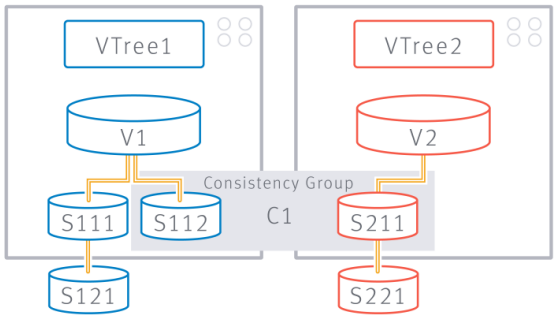

The structure related to all the snapshots resulting from one volume is referred to as a V-Tree (short for volume tree). When taking a snapshot in the system, you can specify more than one volume. All snapshots taken together form a consistency group. They are consistent in the sense that they were all taken at the same time. So if there is a contextual relationship between the data contained by all the snapshot members, then that set is meaningful. The consistency group allows manipulation of the entire set.

If you remove an entire consistency group, all of the snapshots that were taken together will be removed. In Figure 10 on page 62, in RED, S211 is a snapshot of V2. Since S112 and S211 were taken together, they compose a consistency group designated as C1.

Note

The consistency group is only for convenience purposes. There are no protection measures done by VxFlex OS to conserve the consistency group. For example, you can remove a snapshot that is a member of a consistency group.

V-Trees

A V-Tree (short for volume tree) is the structure comprised of a volume and the snapshots resulting from that volume. It is a tree spanning from the source volume at its root, whose descendants are either snapshots of the volume itself or snapshots of a snapshot. In the V-Tree diagram, S_{111} and S_{112} are snapshots of V_1. S_{121} is a snapshot of snapshot S_{111}. Together, V_1 and S_{1xy} are the V-Tree of V_1.

Figure 11 V-Tree diagram

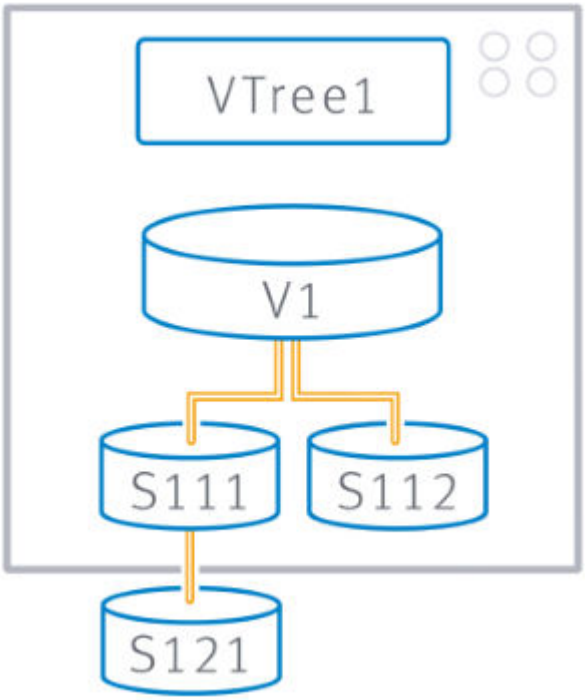

Other functions

VxFlex OS includes the following functions:

- Log Collection
  This operation collects logs from the MDM, SDS, AMS, iDRAC, RAID controller, and OS (including ESXi VM-support) logs. Logs are collected in ZIP format, which you can download for viewing.

  You can collect logs using the VxFlex OS GUI, the AMS CLI, and the REST API.

  The log collection process has the following stages:

  1. Generate the logs on the local nodes.
  2. Transfer the log files to the AMS server.
  3. Download the logs.

  Log collection is described in the *VxFlex OS Log Collection Technical Notes.*
- Upgrade system components, as described in the *VxFlex Ready Node AMS Upgrade Guide*.
  The non-disruptive operation can upgrade the following components:
  - VxFlex OS components (MDM and TieBreaker, SDS, SDC, LIA, and RFcache)
  - DAS Cache
  - Operating system patches
  - VxFlex OS AMS and VxFlex OS GUI
  - BMC (iDRAC)

- Node firmware (Storage controller, NIC, disk)
- BIOS

All upgrades (including OS updates) are performed from the VxFlex OS GUI. Upgrade files are provided via a TAR file, which also includes an XML file that orchestrates the upgrade process.

Note

All VxFlex OS GUI software and hypervisor upgrades (including security patches) in the AMS environment must be orchestrated by AMS. Upgrades performed using any method other than AMS may impact the AMS management capabilities in general, and orchestrated upgrades in particular.

Follow VxFlex Ready Node AMS documentation for available upgrade paths. Should a required path be missing, contact Customer Support.

- AMS management takeover of non-managed nodes
  The AMS can "take ownership" of all the nodes on an existing VxFlex Ready Node system that is currently running VxFlex OS GUI without AMS management. When the process is complete, all the servers in the system are managed with the AMS.

  The take-ownership process is performed from the VxFlex OS GUI together with an AMS CLI command.

Implementing VxFlex OS

Implementing the VxFlex Ready Node system includes the following steps:

1. Assemble the hardware in the rack, and connect to the fully-protected source power supply and a fully HA networking solution.
2. Power on the chassis.
3. Install the AMS and the VxFlex OS GUI, on a server that will remain outside of the VxFlex Ready Node system.
4. Configure the initial IP address pools, passwords and vCenter (for ESXi systems).
5. Run the add node/deployment process.

Communication is done over the existing LAN data network using standard TCP/IP. The MDM and SDS nodes can be assigned up to eight IP addresses, enabling wider bandwidth and better I/O performance and redundancy.

You perform the physical layer setup with the Deployment wizard.

After completing this installation, the physical layer is ready, and can expose a virtual storage layer.

Implementing VxFlex OS in an ESXi-based system

Implementation

In the VMware environment, the MDM and SDS components are installed on a dedicated SVM, whereas the SDC is installed directly on the ESX host.

This implementation is illustrated in the following figure:

Figure 12 VxFlex OS implementation on ESX

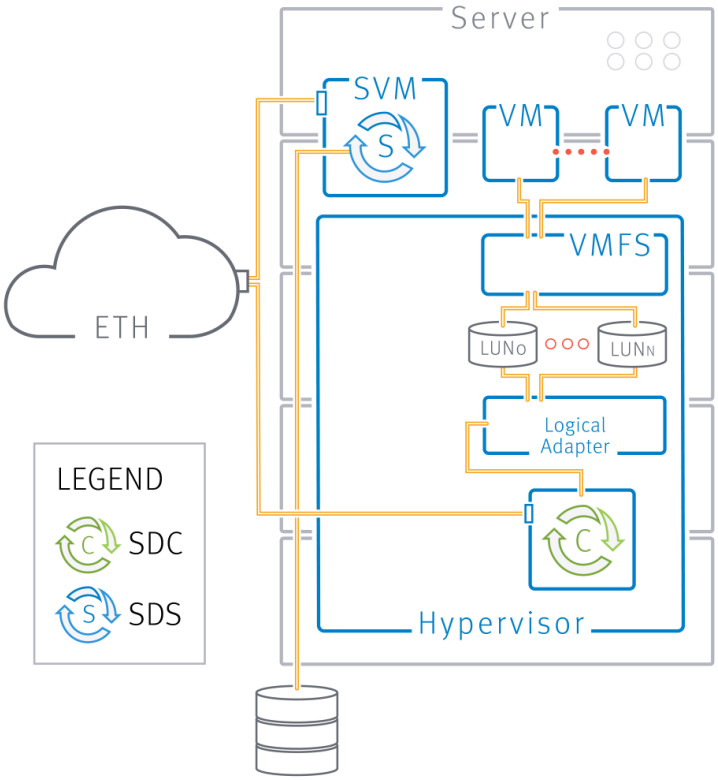

The LUNs in the previous figure can be formatted with VMFS, and then exposed using the ESXi host to the virtual machine, or can be used as RDM devices. When the LUNs are used as RDM devices, the VMFS layer is omitted.

Device management

Devices can be managed in the following ways:

- VMDirectPath I/O
  Device management is performed via the SVM, yielding the best I/O performance. Devices are added to the system after the deployment. On factory-installed VxFlex Ready Node servers, the AMS configures DirectPath on the ESX servers and adds devices to the system.
- RDM
  PowerEdge 14G servers with any NVME devices cannot be added in a DirectPath-based system. AMS uses RDM for NVMe devices. Using RDM mapping, a device is created on the SVM that points to the physical disk on the ESXi node.

Pre-deployment considerations

You should take these considerations into account before deploying the system:

- Are there flash devices that will be added to the SDS?
- Do you want to create Fault Sets? See the requirements in Fault Sets on page 33.

Post-deployment considerations

You should take these considerations into account after deploying the system:

- After deployment is complete, AMS sets all SVMs to start automatically with the system. Do not set SVMs under the VMware resource-pool feature.

Implementing VxFlex OS over a Linux bare metal system

You can implement VxFlex OS on bare metal Linux servers in the following configurations:

- Converged - the MDM, SDS, SDC, and LIA components are installed on the host.
- 2-layer - the MDM, SDS, and LIA components are installed on the host. After deployment, manually install the SDC component on application nodes that consume VxFlex OS storage. This is the default architecture.

Linux VxFlex OS converged implementation

In the Linux VxFlex OS converged environment, the MDM, SDS, and SDC components are all installed on the host.

The LUNs can be formatted with the operating system-supported file system, and then exposed using mount points to the OS, or they can be used as raw devices. When the LUNs are used as raw devices, the file system layer is omitted.

If multiple nodes are accessing the same device a cluster file system solution must be implemented on the client side.

SDS device management
Devices are managed in the following way:

- Direct I/O
  Device management is performed on the host directly, yielding the best I/O performance. Devices are added to the system after the deployment. On factory-installed VxFlex Ready Node servers, the AMS configures the optimal device configuration on the Linux servers and adds devices to the system.

Deployment considerations
Take the following into account:

- Pre-deployment
  - Are there flash devices that will be added to the SDS?
  - Do you want to create Fault Sets? See the requirements in the "Fault sets" section.
- Post-deployment
  - After deployment is complete, AMS sets all VxFlex OS processes to start automatically with boot of the nodes. Do not set change the settings unless instructed by VxFlex OS guidelines.

Linux 2-layer VxFlex OS implementation

In the Linux 2-layer architecture, the MDM and SDS components are all installed on the host. The SDC component is manually installed on the compute node that consumes VxFlex OS storage. This is the default implementation architecture.

The LUNs can be either formatted with the compute operating system-supported file system, and then exposed using mount points to the operating system, or can be used as raw devices. When the LUNs are used as raw devices, the file system layer is omitted.

If multiple nodes are accessing the same device a cluster file system solution must be implemented on the client side.

SDS device management

Devices are managed in the following way:

- Direct I/O
  Device management is performed on the host directly, yielding the best I/O performance. Devices are added to the system after the deployment. On factory-installed VxFlex Ready Node servers, the AMS configures the optimal device configuration on the Linux servers and adds devices to the system.

Deployment considerations

Take the following into account:

- Pre-deployment
  - Are there flash devices that will be added to the SDS?
  - Do you want to create Fault Sets? See the requirements in the "Fault sets" section.
- Post-deployment
  - AMS sets all VxFlex OS processes to start automatically with boot of the nodes. Do not set change the settings unless instructed by VxFlex OS guidelines.
  - Manually install SDC on the application nodes.

SAN virtualization layer

The MDM cluster manages the entire system. It aggregates the entire storage exposed to it by all the SDSs to generate a virtual layer - virtual SAN storage. Volumes can now be defined over the Storage Pools and can be exposed to the applications as a local storage device using the SDCs.

To expose the virtual SAN devices to your servers (the ones on which you installed and configured SDCs), perform the following:

- Define volumes. Each volume defined over a Storage Pool is evenly distributed over all members using a RAID protection scheme. By having all SDS members of the Storage Pool participate, VxFlex OS ensures:
  - Highest and most stable and consistent performance possible
  - Rapid recovery and redistribution of data
  - Massive IOPS and throughput
    You can define volumes as follows:
    - Thick
      Capacity is allocated immediately, even if not actually used. This can cause capacity to be allocated, but never used, leading to wasted capacity.

      Thick capacity provisioning is limited to available capacity.
    - Thin
      Capacity is “on reserve,” but not allocated until actually used. This policy enables more flexibility in provisioning.

      Whereas thick capacity is limited to available capacity, thin capacity provisioning can be oversubscribed, as follows:

      Maximum thin capacity provisioning = 5 * (gross capacity - used capacity)

When capacity usage reaches the level where it may cause IO errors, alerts are generated. At certain higher capacity levels, volumes (even thin volumes) can no longer be created.

Example:

In a system with 3 SDSs, each with 10 TB, there are 30 TB of storage.

In the system, there is already a thick-provisioned volume that takes up 15 TB of the gross capacity (created by adding a 7.5 TB volume).

MDM will allow a total of 300 TB gross to be provisioned, and since 15 TB are already allocated, you can add a thin-provisioned volume of 285 TB gross (by adding a 142.5 TB volume) or a thick-provisioned volume of 15 TB gross.

- Map volumes. Designate which SDCs can access the given volumes. This gives rise to the following:
  - Access control per volume exposed
  - Shared nothing or shared everything volumes

    Once an SDC is mapped to a volume, it immediately gets access to the volume and exposes it locally to the applications as a standard block device. These block devices appear as `/dev/sciniX` where *X* is a letter, starting from "a."

    For example:

    - `/dev/scinia`
    - `/dev/scinib`
  - The maximum amount of partitions for the scini disk is 15.
- The maximum amount of volumes that can be mapped to an SDC is listed in the "Product limits" table.

Note

SDC mapping is similar to LUN mapping, in the sense that it only allows volume access to clients that were explicitly mapped to the volume.

This is the end of the system setup.

Compute cluster configuration and workload deployment

The VxFlex OS GUI and the REST API can be used to administer and orchestrate Compute cluster creation and allocation of nodes. You can use these interfaces to create compute clusters on ESX servers, configure the application network and ports, map storage, and then deploy Virtual Machine workloads. This feature also lets you monitor and modify clusters and workloads that were created using the VxFlex OS GUI or REST API.

This feature simplifies configuration of the application network and ports, and provides an interface for monitoring, editing, and removing compute nodes and clusters that have been configured using the VxFlex OS GUI or REST API.

Volume creation and ESX mapping processes include the ability to format the mapped volume with a VMFS datastore on ESX with the VMFS5 format.

Cluster and network configuration can be performed in simple or advanced modes. For some settings, auto-select options allow the system to recommend the most suitable configurations. In addition, default or customized profiles speed up the configuration activities.

During the provisioning process, overcommit can be enabled. VM allocation in VMware is based on cores and RAM, and it is not recommended to exceed either of them. However, due to virtualization and real use cases, especially CPU usage, the core allocation may exceed the actual cores available.

The total number of cores in a cluster is:

- Core per Node per socket * 2 because of hyper threading
- Number of sockets in Hyper Converged = 2
- SVMs (Storage Virtual Machines) = 8 vCPU
- Add up all the cores per node to the total. Example if all nodes are even and there are 3 SVMs:
  Core: 3 nodes * 2 sockets * 16 cores *2 (for HT) - 24 (3 SVMs) = 168

When a template is selected, you can decide the number of cores and RAM with which it will deploy the cluster.

Limitations in this version:

- The supported versions of VMware Tools that are compatible with the customization of workload VMs are: 9.4.15 - 10.1.7
- A node added manually to a cluster (not via the VxFlex OS GUI) will not be added to the compute environment managed by the VxFlex OS GUI.
- IPv6 is not supported by this feature.
- A workload cannot be expanded if the template has more than one VM network configured with the same portgroup.

Maintenance

Maintenance of VxFlex OS is primarily limited to configuration changes of the physical and virtual layers. It requires minimal user attention. When maintenance or planned restart of an SDS is required, the maintenance mode feature can be used to streamline system operation.

Maintaining the physical layer

In the physical layer, maintenance is limited to adding and removing hardware units and configuring them into VxFlex OS. These operations are usually a result of:

- Scaling out when there is a need for additional capacity. This usually results in adding more storage media to the existing servers, or adding additional servers.
- Hardware failure. In cases where there is a hardware (storage media or server) failure and it needs to be replaced.

  In all of the above cases, the operation will require adding or removing storage capacity from the system. In some cases, it may include adding or removing an entire server, and its associated storage media, from the configuration. As far as VxFlex OS is concerned, all of these activities translate to SDS reconfigurations.

  If the removed server is an SDC node, or the server to be added requires exposing storage locally, SDC reconfiguration will happen as well.
- Adding or removing storage media. Add or remove the media from the SDS with which it is associated. VxFlex OS will redistribute the data accordingly and seamlessly.

Instant maintenance mode

Instant maintenance mode enables you to restart a server that hosts an SDS, with minimal impact on VxFlex OS, thus bypassing the disruption and effort caused by disorderly shutdown, Protection Domain shutdown, and orderly shutdown.

Whereas VxFlex OS always uses two copies of user data, invoking maintenance mode introduces an additional copy that stores all writes created during maintenance to an SDS or Fault Set (created during maintenance) in both a primary location and a new location. This copy prevents data loss if a single failure occurs.

When the SDS or Fault Set is returned from maintenance mode, only the new writes are required to be resynchronized, thus minimizing data transfer during and after the update.

Instant maintenance mode does not interrupt application IOs; it can be run on any amount of members of a Fault Set; and it can run in parallel on different Protection Domains. While an SDS is in maintenance mode, most VxFlex OS operations (like adding a volume) cannot be performed in the Fault Set, Protection Domain, or Storage Pool in which the SDS and its devices reside.

To invoke maintenance mode, the following conditions are required:

- Only one Fault Unit (or standalone SDS) can be in maintenance mode at any given time.
- No other SDSs can be in degraded or failed state (force override can be used).
- There must be adequate space on other SDSs for the additional backup (force override can be used).

  Note

  Use of force override options when entering maintenance mode can lead to data unavailability while maintenance mode is activated.

  While an SDS is in maintenance mode, it can be shut down with no danger to data.

Adding or removing a server

When adding or removing servers, VxFlex OS redistributes the data accordingly and seamlessly.

Maintaining the virtualization layer

The following operations may be performed on volumes that are exposed by the VxFlex OS virtual SAN:

- Add or remove a volume:

  Create or delete a volume in the system.
- Increase volume size:

  Add capacity to a given volume, as needed. The change in volume size occurs seamlessly without interrupting I/O.
- Map and unmap volumes to an SDC:

  This enables or disables access to a volume by an SDC, and thus by an application residing on the same node.

Management tools

You can provision, maintain, and monitor VxFlex OS using the following management clients:

- Command Line Interface (CLI)

  The CLI enables you to perform the entire set of configure, maintain, and monitor activities in VxFlex OS.

- Graphical User Interface (GUI)

  The VxFlex OS GUI enables you to perform standard configure and maintain activities, as well as to monitor the storage system's health and performance. You can use the VxFlex OS GUI to view the entire system, and then drill down to different elements.

- REST gateway

  A REST API can be used to expose monitoring and provisioning via the REST interface. The REST server is installed as part of the VxFlex OS AMS.

  Many VxFlex OS activities can be performed in more than one management tool.

Firmware patching

During the deployment or extend processes, hardware components are checked automatically for the appropriate firmware versions. If mismatches are found, the firmware is upgraded or downgraded to the correct version.

Firmware patching of hardware components primarily happens during an orchestrated non-disruptive upgrade process. On ESXi-based systems, firmware patching is not a service-affecting operation, because ESXi servers can operate in maintenance mode. However, in bare-metal scenarios where SDCs are affected, system restarts may be required. In such cases, user confirmation will be required.

Each VxFlex Ready Node version has its own set of supported drivers and firmware. The following tables summarize the drivers and firmware supported by this version.

Table 21 13G PowerEdge Servers

Item	Firmware Version	Additional Details
Node		
BIOS	2.7.1	
BMC	2.52.52.52 A00	
Backplane	2.25 A00	
Expander BP	3.35	
OS Driver Pack	17.08.10 A00	
H730 RAID controller firmware	25.5.4.0006 A12	PERC H730\H730p RAID controller
SAS controller HBA330 Mini firmware	15.17.09.06	HBA330 SAS controller

Table 21 13G PowerEdge Servers (continued)

Item	Firmware Version	Additional Details
NIC 1G FW	17.5.10	I350 Gigabit Network Connection
NIC 10G X520/540 FW	17.5.10	82599 10 Gigabit Dual Port Network Connection
NIC 10G X710 FW	17.5.10	Follow instructions to set the network card - TSO, LRO - OFF
Storage Disks		
Seagate 1.2TB ST1200MM0088 WXPCX	TT31	
HGST 1.2TB HUC101812CSS204 9XNF6	FJ34	
Toshiba PM3 RI 4KG4X, R87FK, M09K5	AM07	
Toshiba PM3 MU M91TJ, 77K16	AM07	
Toshiba PM3 MU3 YYC10, 4XC39, GYMY9	AM07	
HGST HUSMM WI CW988, 5HR3X	L384	
Samsung RI PM1633a ISE 7FNRX 086DD JR1HP	DSL7	
Toshiba PM4 RI ISE MWGK7 0FYFW XCRDV	AS0B	
Toshiba PM4 MU3 ISE 503M7 V0K7V 3DDFT	AS0B	
Toshiba PM4 WI ISE CN3JH, GVTYD, R1YFC	AS0B	
Seagate 1.2TB ISE ST1200MM0198 FR6W6	UT71	
HGST 1.2TB ISE HUC101812CSS200 0KV02	FU29	
Samsung PM863a RI ISE T2G0Y 9W12R 9Y3HD	GC57	
Seagate Skybolt 1.2TB G2G54	ST31	

Table 22 14G PowerEdge Servers

Item	Firmware Version	Additional Details
Node		

Table 22 14G PowerEdge Servers (continued)

Item	Firmware Version	Additional Details
BIOS	1.3.7	
BOSS Controller FW	2.5.13.2008 A00	
M2 Disk FW	N201DL42	
Backplane	4.23 A00	
Expander BP	2.17	
iDRAC/BMC	3.15.15.15 A00	
OS driver Pack	17.10.13 A00	
RAID controller H730p firmware	25.5.4.0006 A12	PERC H730\H730p RAID controller
SAS controller HBA330 firmware	15.17.09.06	SAS controller HBA330 SAS controller HBA330 driver
NIC 1G FW	18.3.6 A00	X350 Gigabit Network Connection NIC 1G Driver
NIC 10G Mellanox CX-4 FW	14.20.18.20 A00	
NIC 10G X710 FW	18.3.6 A00	In Linux, follow doc instructions to set card network - TSO, LRO - OFF
NIC 10G X550 FW	18.3.6	X350 10 Gigabit Base T Network Connection
NIC 10G X550 Driver	4.4.1-1OEM.550.0.0.1331820	X350 10 Gigabit Base T Network Connection
	4.4.1-1OEM.550.0.0.1331820	
	4.4.0-k-rh7.3/5.2.4	
	4.2.1-k	
	3.14.3.0 / 4.1.42.0	
iSM	3.0.1	Cells on the left relate to the corresponding cells in the row above.
	3.0.1	
	3.0.1	
	3.0.1	
	3.0.1	
	TBD	
Storage Disks		
Seagate 1.2TB ISE ST1200MM0198 FR6W6	UT71	
Seagate Skybolt 1.2TB G2G54	ST31	

Table 22 14G PowerEdge Servers (continued)

Item	Firmware Version	Additional Details
HGST 1.2TB ISE HUC101812CSS200 0KV02	FU29	
Toshiba 1.2TB ISE AL14SEB120NY 3K30N	EA02	
Intel S3520 VXG5N R3J3Y	DL43	
Samsung MU PM1635a ISE HF06W W5PP5 8RC8K	DWL7	
Samsung 1725a	1.0.4 A00	
Samsung PM863a RI ISE T2G0Y 9W12R 9Y3HD	GC57	
Samsung SM863a MU ISE DD4G0 K5P0T	GD53	
Toshiba PM4 RI ISE MWGK7 0FYFW XCRDV	AS0B	
Toshiba PM4 MU3 ISE 503M7 V0K7V 3DDFT	AS0B	
Intel S3520	DL42	
Samsung RI PM1633a ISE 7FNRX 086DD JR1HP	DSL7	
Toshiba HK4 31Y1M, V7HCN	DAC9	
Seagate Skybolt 2.4TB RWR8F	ST51	

PART 2

Getting Started

The following chapters describe how to get VxFlex OS started in your environment. Chapters include:

Chapter 3, "User Management"

Chapter 4, "Creating and Mapping Volumes"

CHAPTER 3

User Management

The following topics describe how to create and manage users.

- Overview 78
- User roles 78
- Authenticating with the management server 79
- Adding and modifying local users 83

Overview

VxFlex OS supports local domain user authentication, and LDAP domain authentication. In addition, secure authentication is used between system internal and external components. This chapter provides the CLI commands used to create and manage VxFlex OS users. The REST API can also be used to configure LDAP. For more information, see the operations for MDM clusters in the *VxFlex OS REST API Reference Guide*.

- To set up local domain users, follow the instructions in this chapter.
- To set up LDAP users, see a detailed explanation in the document *VxFlex OS User Roles and LDAP Technical Notes*. In general, the following steps must be performed:

1. Add LDAP service to the MDM.
2. Create Active Directory (AD) groups that correspond to the user roles offered by VxFlex OS.
3. Set the system-wide authentication method (use with caution, because it is complex to roll-back this operation).
4. Log in again to apply the changes that you made.

User roles

The authorization permissions of each user role are defined differently for local authentication, and for LDAP authentication. Although the role names are similar, the permissions granted to them are not.

User roles defined in the LDAP domain are mutually exclusive, with no overlap—with the exception of the Configurator role. If you want to give an LDAP user permission to perform both monitoring and configuration roles, for example, assign that user to both the Backend/Frontend Configurator and Monitor LDAP groups.

The Configurator and Super User roles do not exist at all for LDAP.

The following table describes the permissions that can be defined for local domain users and for LDAP domain users.

Table 23 Local and LDAP user roles and permissions

User role	Query		Configure parameters		Configure user credentials	
	Local	LDAP	Local	LDAP	Local	LDAP
Monitor	Yes	Yes	No	No	No	No
Configurator (this role is only applicable for local users)	Yes	Not applicable	Yes (an aggregation of both Frontend and Backend Configurator)	Not applicable	No	Not applicable
Backend Configurator	Yes	No	Yes Backend operations only (Protection Domains,		No	No

Table 23 Local and LDAP user roles and permissions (continued)

User role	Query		Configure parameters		Configure user credentials	
	Local	LDAP	Local	LDAP	Local	LDAP
			Storage Pools, Fault Sets, SDSs, Devices, other system settings)			
Frontend Configurator	Yes	No	Yes Frontend operations only (Volumes, SDCs, Snapshots)		No	No
Administrator	Yes	No	Yes	No	May configure Configurator and Monitor users	
Security Roles	No	No	No	No	May define Administrator users and control LDAP	
Super User (only one Super User is allowed per system, and it must be a local user)	Yes	Not applicable	Yes	Not applicable	Yes	Not applicable

Authenticating with the management server

To access the CLI, you must first log in, using a terminal application. Once logged in, you can set the user authentication method.

If the CLI and the AMS do not reside on the same server, add the `--mdm_ip` parameter, the IP address of the AMS server, to all CLI commands. For example:

```
scli --mdm_ip 10.10.10.3 --login --username admin
```

You will be prompted to enter the password.

For AMS login, the password is created during the installation of the AMS.

When using the LDAP, the command should also include the LDAP domain it is using, for example:

```
scli --mdm_ip 10.10.10.3,10.10.10.4 --login --username
JohnDoe@ldap.acme.com
```

The default user created during setup is the Super User, with the *admin* username.

Note

For detailed information on issuing CLI commands, see the *CLI Reference Guide*.

Logging in

Command

login

Syntax

```
scli --login --username <NAME> [--password <PASSWORD>] [(--
ldap_authentication | --native_authentication)] [--
approve_certificate]
```

Actual command syntax is operating-system dependent. For more information, see "CLI basics".

Description/Notes

This command logs the specified user into the management system. Every user must log in before performing CLI commands.

When a user is authenticated by the system, all commands performed will be done with their respective role until a logout is performed, or until the session expires, by reaching one of the following timeouts:

- Maximum session length (default: 8 hours)
- Session idle time (default: 10 minutes)

  To log in to the AMS, from a local Linux AMS server, run:

  ```
  /opt/emc/scaleio/ams/webapps/ROOT/resources/scli --login --
  username admin --password <AMS_admin_password>
  ```

  Example

  ```
  /opt/emc/scaleio/ams/webapps/ROOT/resources/scli --login --
  username admin --password Scaleio123
  ```

  Output, similar to the following, is displayed:

  ```
  Logged in. User role is Administrator. System ID is
  0000000000000000
  ```

  To log in to the AMS, via SSH to the AMS server, run:

  ```
  /opt/emc/scaleio/ams/webapps/ROOT/resources/scli --login --
  username admin --password <AMS_admin_Password> --mdm_ip <AMS_IP>
  ```

Example

```
/opt/emc/scaleio/ams/webapps/ROOT/resources/scli --login --
username admin --password Scaleio123 --mdm_ip 10.136.247.106
```

Output, similar to the following, is displayed:

```
Logged in. User role is Administrator. System ID is
0000000000000000
```

Parameters

Parameter	Description
--username	User name
--password	User password. If you do not type your password, you will be prompted to do so. **Note** In Linux, to prevent the password from being recorded in the history log, leave out the `--password` flag and enter the password interactively.
--ldap_authentication	Login with LDAP authentication method. LDAP authentication parameters should be configured and LDAP authentication method should be set.
--native_authentication	Login with the native authentication method (default)
--approve_certificate	Preemptive approval of the MDM certificate

Example

```
scli --login --username siouser1 --password 1!2@3A
```

Logging out

Command

`logout`

Syntax

```
scli --logout
```

Actual command syntax is operating-system dependent. For more information, see "CLI basics".

Description/Notes

This command logs the current user out of the system.

Example

```
scli --logout
```

Setting the User Authentication Method

Command

set_user_authentication_method

Syntax

```
scli --set_user_authentication_method (--ldap_authentication | --
native_authentication | --native_and_ldap_authentication) [--
i_am_sure]
```

Actual command syntax is operating-system dependent. For more information, see "CLI basics".

Description/Notes

This command sets user authentication method for the system.

Note

Use this command with caution, because it is complex to roll-back this operation.

Parameters

Parameter	Description
--ldap_authentication	LDAP based authentication method, where users are managed in LDAP compliant server. In order to switch to this authentication method, LDAP service parameters should be configured.
--native_authentication	Native authentication method, where users are managed locally in the system
--native_and_ldap_authentication	A hybrid authentication method. Both LDAP and Native users may log into the system after it is set.
--i_am_sure	Preemptive approval

Example

```
scli --set_user_authentication_method --
native_and_ldap_authentication --i_am_sure
```

Adding and modifying local users

Any user with the administrator role can display all of the users in the system, their roles, and their user ID, and can change the password of a user that is currently logged into the system.

The following CLI commands allow you to manage local users.

Adding users

Command

add_user

Syntax

```
scli --add_user --username <NAME> --user_role <Monitor|Configure|
BackEndConfigure|FrontEndConfigure|Security|Administrator>
```

Actual command syntax is operating-system dependent. For more information, see "CLI basics".

Description/Notes

This command, only available to administrator roles, adds a user to the system. A randomly generated password for the created user is returned.

Note

When adding users, you may need to use the " or ' characters to enable the use of Bash syntax characters.

Parameters

Parameter	Description
--username	The user name to add to the system
--user_role	The user role: • Monitor • Configure • BackEndConfigure • FrontEndConfigure • Security • Administrator

Parameter	Description
	For more information about the roles, see Table 20.

Example

```
scli --add_user --username siouser2 --user_role Configure
```

Deleting users

Command

delete_user

Syntax

```
scli --delete_user (--user_id <ID> | --username <NAME>)
```

Actual command syntax is operating-system dependent. For more information, see "CLI basics".

Description/Notes

This command, only available to administrator roles, deletes the specified user from the system.

Parameters

Parameter	Description
--user_ID <ID>	ID of the user to be deleted
--username <NAME>	User name of the user to be deleted

Example

```
scli --delete_user --username siouser2
```

Modifying user credentials

Command

modify_user

Syntax

```
scli --modify_user (--user_id <ID> | --username <NAME>) --user_role
<Monitor|Configure|BackEndConfigure|FrontEndConfigure|Security|
Administrator>
```

Actual command syntax is operating-system dependent. For more information, see "CLI basics".

Description/Notes

This command, only available to administrator roles, modifies the user role of the specified user in the system.

Parameters

Parameter	Description
--user_id	The ID number of the user to modify. You can obtain the user ID from the `query_users` command, or when adding the user.
--username	The user name of the user to modify
--user_role	The new user role: • Monitor • Configure • BackEndConfigure • FrontEndConfigure • Security • Administrator For more information about the roles, see Table 20.

Example

```
scli --modify_user --username siouser3 --user_role Monitor
```

Displaying system users and their roles

Command

`query_users`

Syntax

```
scli --query_users
```

Actual command syntax is operating-system dependent. For more information, see "CLI basics".

Description/Notes

This command, only available to administrator roles, displays all the users defined in the system, their roles, and user ID.

Example

```
scli --query_users
```

Displaying information about a specific user

Command

query_user

Syntax

```
scli --query_user (--user_id <ID> | --username <NAME>)
```

Actual command syntax is operating-system dependent. For more information, see "CLI basics".

Description/Notes

This command, only available to administrator roles, displays information about the specified user.

Parameters

Parameter	Description
--user_id	The user ID of the user. You can obtain the user ID via the `query_users` command, or when adding the user.
--username	The name of the user

Example

```
scli --query_user --username sio_user
```

Resetting user passwords

Command

reset_password

Syntax

```
scli --reset_password (--user_id <ID> | --username <NAME>)
```

Actual command syntax is operating-system dependent. For more information, see "CLI basics".

Description/Notes

This command, only available to administrator roles, generates a new password for the specified user. The user will need to change the password again after logging in with this password.

Parameters

Parameter	Description
--user_id	The user ID of the user whose password will be reset. You can obtain the user ID from the `query_users` command, or when adding the user.
--username	The user name of the user whose password will be reset

Example

```
scli --reset_password --username siouser3
```

Changing user passwords

Command

`set_password`

Syntax

```
scli --set_password [--old_password <OLD_PASSWORD>] [--new_password
<NEW_PASSWORD>]
```

Actual command syntax is operating-system dependent. For more information, see "CLI basics".

Description/Notes

This command changes the password of the user that is currently logged in to the system.

Parameters

Parameter	Description
--old_password	User's current password. **Note** In Linux, to prevent the password from being recorded in the history log, leave out the `--old_password` flag and enter the password interactively.
--new_password	User's new password. **Note** In Linux, to prevent the password from being recorded in the history log, leave out the `--new_password` flag and enter the password interactively. Password rules: • Must be between 6 (default, configurable) and 31 characters • Presence of least 3 groups out of [a-z], [A-Z], [0-9], special characters (!@#$ ...) • The current password is not allowed

Example

```
scli --set_password --old_password 1!2@3A --new_password P9*7&6
```

Disabling the Default Admin User

Command

`disable_admin`

Syntax

```
scli --disable_admin [--i_am_sure]
```

Actual command syntax is operating-system dependent. For more information, see "CLI basics".

Description/Notes

This command disables the default Super User. The Super User is a new role that replaces the default administrator in the earlier versions. This is the default user for

setting up the system. It has all the privileges of all the user roles. The Super User is only available in the local domain, and cannot be an LDAP user. One may disable the Super User in order to ensure that all users are associated with specific people. The Super User role should be used in two cases:

- When setting up a new system
- When there is no desire to have separate users and there is no desire to define additional users.

  The default credentials are:

  Username: `admin`

  Password: `admin`

  Note

  Having a single Super User dictates that the audit log will not track a certain person per se.

  If you need to re-enable the Super User, use the `reset_admin` command.

Parameters

Parameter	Description
--i_am_sure	Preemptive approval

Example

```
scli --disable_admin --i_am_sure
```

CHAPTER 4

Creating and Mapping Volumes

The following topics describe how to create volumes from devices added to SDS nodes, and then to map the volumes to SDC nodes. Devices may have been added during, or after, the installation process.

- Creating and mapping volumes overview 92
- Creating volumes 92
- Mapping a volume to an SDC 94

Creating and mapping volumes overview

You can create volumes from devices added to SDS nodes, and then map the volumes to SDC nodes. Devices may have been added during, or after, the installation process.

Creating and mapping volumes can be performed using various management tools. The following procedures describe how to do so with the VxFlex OS GUI, and provide references to the other tools.

VxFlex Ready Node AMS users should always perform these activities using the VxFlex OS GUI.

The creating and adding volume process described in this section is necessary, as part of the Getting Started process, before applications can access the volumes. In addition, you may create additional volumes and map them as part of the maintenance of the virtualization layer.

Creating volumes

Adding volumes

Add volumes to a system.

Before you begin

There must be at least three SDS nodes in the system and there must be sufficient capacity available.

Note

For the minimum size of an SDS, see System requirements on page 21.

The adding and mapping volume process is necessary, as part of the getting started process, before applications can access the volumes. In addition, you may create additional volumes and map them as part of the maintenance of the virtualization layer.

You can configure the caching option when creating the volumes, or you can change the Read RAM Cache feature later. If you want to enable the caching feature, ensure that the feature is also enabled in the backend of the system, for the corresponding Storage Pool and SDSs. For more information, see Changing Read RAM Cache volume settings on page 193.

Define volume names according to the following rules:

- Contains less than 32 characters
- Contains only alphanumeric and punctuation characters
- Is unique within the object type

VxFlex OS objects are assigned a unique ID that can be used to identify the object in CLI commands. You can retrieve the ID via a query, or via the object's property sheet in the VxFlex OS GUI. It is highly recommended to give each volume a meaningful name associated with its operational role.

To add one or multiple volumes, perform these steps:

Procedure

1. In any of the **Frontend** > **Volumes** views, navigate to the Storage Pool to which you want to add the volume, and select it.
2. From the **Command menu** or context-sensitive menu, select **Add Volume**.
3. In the **Add Volume** window, if you want to create more than one volume, select **Create multiple volumes** and type the number of volumes you would like to add in the **Copies** box.
   - If you type `1`, only one volume will be created (optional—can be left blank).
   - If you type a number greater than 1, the characters `%i%` will be added to the **Name** box, and multiple volumes will be created, accordingly.
     The volumes will be named and numbered automatically, starting from 1. If you want the numbering to start from a different number, type it in the **Start numbering at** box, as described in Step 5. The remaining options in the window will be assigned to all the volumes created in this operation.
4. Type a name for the volume:
   - If you are adding one volume, enter the name in the **Name** box.
   - If you are adding multiple volumes, enter the base name in the **Base name** box.
     The volumes will all be created with the same name, and a number will be appended instead of the characters `%i%`. These characters can be positioned anywhere in the name. The names that will be created are displayed in the right pane of the window, as shown in the figure later in this topic.
5. If you want the numbering to start from a specific number other than 1, type it in the **Start numbering at** box.

   This number will be the first number in the series that will be appended to the volume name. For example, if the **Name** is `Vol%i%` and the **Start numbering at** value is `100`, the name of the first volume created will be Vol100, and the second volume will be Vol101, and so on.
6. Type a number in the **Size** box, representing the volume size in GB (basic allocation granularity is 8 GB).
7. Select either **Thick** (default) or **Thin** provisioning options.
8. If you want to enable the RMcache feature (disabled by default), select **Use RMcache**.
9. Click **OK**.

   The progress of the operation is displayed at the bottom of the window. It is recommended to keep the window open until the operation is completed, and until you can see the result of the operation.

Figure 13 Add Volume window

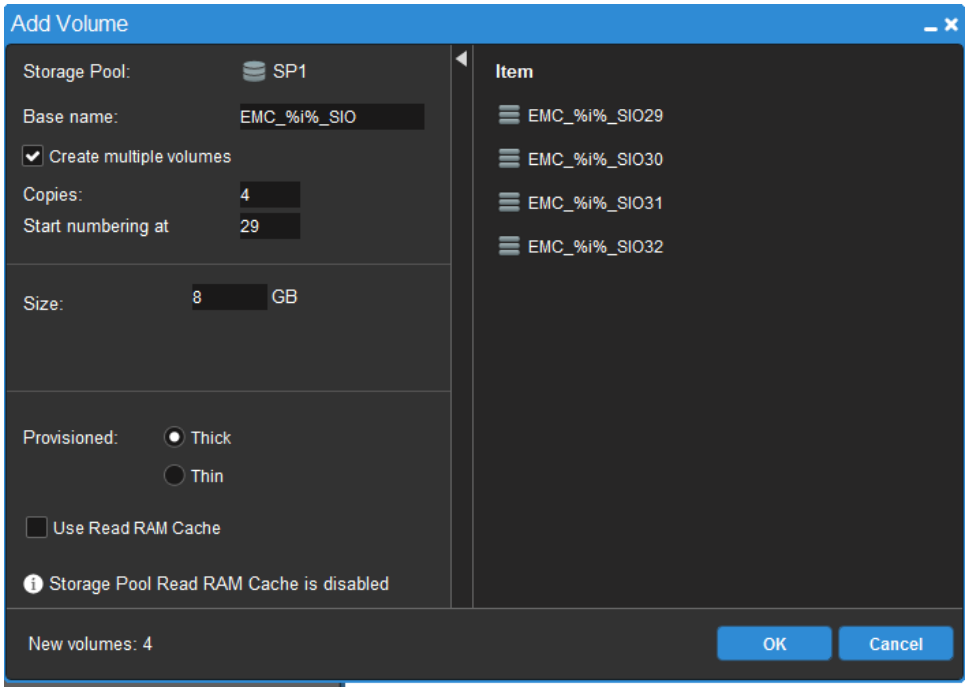

After you finish

To use the created volume, you must map it to (at least) one SDC. If the restricted SDC mode is enabled for the system, you must approve SDCs prior to mapping volumes to them. For more information on approving SDCs, see Approve SDCs using the VxFlex OS GUI on page 199. For more information on mapping volumes, see Mapping a volume to an SDC on page 94.

Mapping a volume to an SDC

Mapping volumes

Map one or more volumes to SDCs.

Mapping exposes the volume to the specified SDC, effectively creating a block device on the SDC.

For Linux devices, the `scini` device name can change on reboot. It is recommended to mount a mapped volume to the VxFlex Ready Node unique ID, a persistent device name, rather than to the `scini` device name.

To identify the unique ID, run the `ls -l /dev/disk/by-id/` command. For more information, see Associating VxFlex OS volumes with physical disks on page 248. You can also identify the unique ID using VMware. In the VMware management interface, the device is called **EMC Fibre Channel Disk**, followed by an ID number starting with the prefix **eui**.

To map volumes, perform these steps:

Procedure

1. In the **Frontend** > **Volumes** view, navigate to the volumes, and select them.
2. From the **Command menu** or context-sensitive menu, select **Map Volumes**.

The **Map Volumes** window is displayed, showing a list of the volumes that will be mapped.

3. In the **Select Nodes** panel, select one or more SDCs to which you want to map the volumes.
   - You can use the search box to find SDCs.
   - If you select an SDC that is already mapped to the volume, a green icon will appear in the mapping matrix on the right side of the window.

4. Click **Map Volumes**.

   The progress of the operation is displayed at the bottom of the window. It is recommended to keep the window open until the operation is completed, and until you can see the result of the operation.

Figure 14 Map Volumes window after mapping is complete

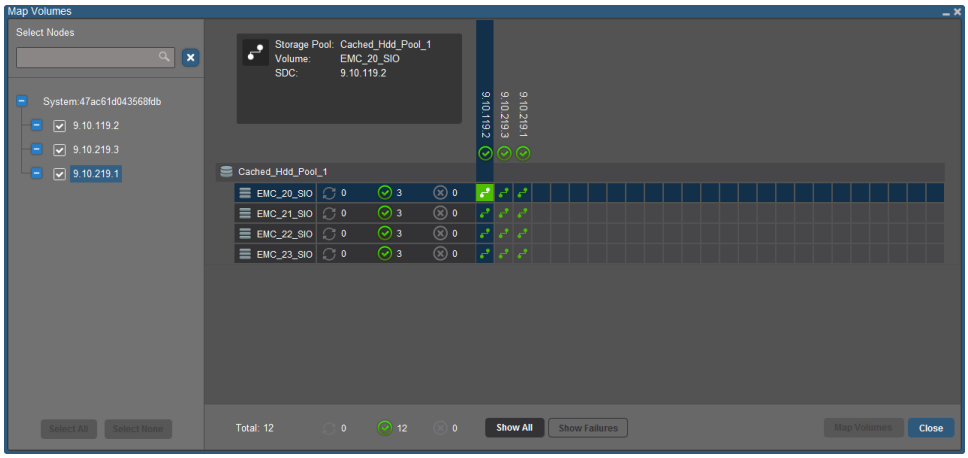

PART 3

Managing and Monitoring

By the time you have reached this part of the guide, VxFlex OS should be up and running, all the way through mapping volumes to SDC nodes and installing the license. The topics in this part of the guide describe how to use the CLI to manage and monitor VxFlex OS activity and components.

Chapters include:

Chapter 5, "Managing System Objects"

Chapter 6, "Security Management"

Chapter 7, "Opening the VxFlex OS GUI and Logging In"

Chapter 8, "VxFlex OS GUI Features"

Chapter 9, "Monitoring the System using the VxFlex OS GUI"

Chapter 10, "Configuring the System using the VxFlex OS GUI"

CHAPTER 5

Managing System Objects

The following topics describe how to manage and configure VxFlex Ready Node system objects.

- CLI basics 100
- Syntax 101
- Managing the VxFlex OS system 103
- Managing the MDM cluster 103
- Managing the SDSs and cache 113
- SDC operations 116
- Managing ESX servers 124

CLI basics

The VxFlex OS CLI, SCLI, enables you to perform all provision, maintain, and monitor activities.

SCLI is installed as part of the AMS, and can be found in the following path:

- Windows:

  `C:\Program Files\EMC\ScaleIO\AMS\webapps\ROOT\resources`

- Linux:

  `/opt/emc/scaleio/ams/webapps/ROOT/resources/`

  Before using the CLI in Linux, you must give the file execute permissions:

```
chmod a+x /opt/emc/scaleio/ams/webapps/ROOT/resources/scli
```

Before using SCLI, you must log in to the AMS using the `scli --login` command, as described in Authenticating with the management server on page 79.

All SCLI commands use one of the following formats, depending on your environment:

NOTICE

These instructions apply to every SCLI command, with one exception: When first logging in using the `--login` command, do not use the `--mdm_port` flag.

- To run commands locally on the AMS, from the correct path enter: `scli <command> --mdm_port 6611` (for --login, do not use --mdm_port)
- To run commands from a Linux server to a remote AMS, from location of the SCLI bin on the local server enter:

```
scli <command> --mdm_ip <AMS_IP_Address> --mdm_port 6611
```

- To run commands from Windows to a remote AMS, from the location of SCLI bin on the local server enter:

```
scli.exe <command> --mdm_ip <AMS_IP_address> --mdm_port 6611
```

Description: Execute a CLI command.

Parameter	Description
--mdm_ip <IP>	Use the IP address of the AMS server.
--approve_certificate	Preemptive approval of the MDM's certificate
<command>	Command to be executed
--mdm_port <PORT>	The port of the MDM that is listening for the command requests. This value is always 6611.

`scli --mdm_ip 10.10.10.3,10.10.10.4 --query_all --mdm_port 6611`

The `mdm_ip` indicates the MDM that receives and is to execute the command. If the command is run from the Master MDM, this switch may be omitted.

To avoid using the `--mdm_ip` parameter in every command, or avoid having to install the CLI on other servers, use SSH or RDM to log in to the shell running on the management server.

Note

- The order of the parameters and command is insignificant.
- SCLI commands are lowercase and case-sensitive.
- All parameters are preceded by `--`

Before using most SCLI commands, you must log in, as described in the *CLI Reference Guide*.

Using SCLI in non-secure mode

If VxFlex OS is running in non-secure mode, you must disable secure communications on every MDM server to enable execution of commands.

- To disable secure communications in Windows, on each MDM open the SCLI `conf.txt` file, and add the following line:

```
cli_use_secure_communication=0
```

- To disable secure communications in Linux, run the following on each MDM:

```
echo cli_use_secure_communication=0 >> ~/.scli/conf.txt
```

Note

For more information on how to set up secure or non-secure mode, see the *VxFlex OS User Guide* .

Syntax

All names of objects in the system will be capitalized, for example, Protection Domain. In the case where the name is in fact initials it will be in uppercase, for example, MDM.

The actual CLI command format uses the following format:

- Message - Required
- <> - Argument
- () – Required element
- [] – Optional element
- | - Select from options A|B|C

  Each command entry uses the above syntax and looks like the following example:

  Usage:

```
scli --cmd_example --r1 (--r2 | --r3 <V1>) [o1 <V2>|o2]
[Options]
```

  Description: a description of what `cmd_example` does

Parameter	Description
--r1	r1 description
--r2	r2 description
--r3 <V1>	r3 description with possible V1 input values
--o1 <V2>	o1 description with possible V2 input values
--o2	o2 description
Options: CHOOSE SEVERAL	
--so1 <V3>	so1 description with possible V3 input values
--so2	so2 description

The interpretation is as follows:

- The text `scli --cmd_example r1` is mandatory.
- `(--r2 | --r3 <V1>)` indicates that you must choose one of the options separated by "|". Selecting an option is REQUIRED, indicated by "( )".
- `[o1 <V2>|o2]` indicates that you may choose one of the options separated by "|". Selecting an option is OPTIONAL, indicated by "[ ]".
- `[Options]` indicates that you may choose one of the options that will be described in the table under Options. It is OPTIONAL, indicated by "[ ]".

Getting help with the CLI

The CLI supports auto-completion. To complete a command or parameters, press the TAB key while typing CLI commands.

Alternately, you can run the `help` command.

Command

`help`

Syntax

```
scli --help [Options]
```

Description/Notes

Use this command to view CLI help.

Parameters

Parameter	Description
Options: CHOOSE ONE	
`mdm`	Help on MDM commands
`sds`	Help on SDS commands

Parameter	Description
`vol`	Help on volume-related commands
`sdc`	Help on SDC commands
`general`	Help on general commands
`all`	Help on all commands
`user`	Help on commands related to user management
`<blank>`	List options

Example

```
scli --help --mdm
```

Managing the VxFlex OS system

This section describes how to manage the VxFlex OS system.

Extending an existing VxFlex OS system

Options for adding nodes to an existing system.

You can add nodes to an existing system, as well as extend the MDM cluster from a 3-node to a 5-node cluster. Depending on your system, you can use the VxFlex OS Installation Manager (for physical servers) or the vSphere plug-in (for ESXi servers).

These topics are described in the sections of the *VxFlex OS Deployment Guide*:

- "Extend an existing VxFlex OS system"
- "Extend the MDM cluster from 3 to 5-node"

Managing the MDM cluster

You can replace or update the server IP address of an MDM that is currently a member of the MDM cluster. Various management tools are available to configure virtual IP addresses for the MDMs.

You can use the VxFlex OS GUI to transfer an MDM cluster component or remove a standby MDM. You can also rename an MDM cluster node.

Replacing, or updating an IP address on a member of the MDM cluster

This section describes how to replace an MDM server that is a member of the MDM cluster. This could be necessitated by a need to replace a faulty server or a need to change the server IP address. For purposes of this section, we will refer to the server that needs to be replaced as the current server.

There are two ways to accomplish this task, determined by whether you are able to add a new server to the cluster, or whether you do not have an extra server to add. If you have an extra server to replace the current server, then there is no need to change the cluster mode (3-node or 5-node). However, if you do not have an

additional server, you will need to reduce the cluster mode, from 5-node to 3-node, or from 3-node to single node.

Note

It is not recommended to use single mode in production systems, except in temporary situations.

Regardless of the circumstances, the following rules are true:

- To remove a cluster member, you first make it a standby, then remove the standby.

  To add a member to a cluster, you first make it a standby, then add the standby to the cluster. In other words, you cannot move a server from being a cluster member to being entirely external, in either direction, without being a standby first.
- The cluster must always have 5, 3, or 1 members, never any other amount.

  For a further understanding of this subject, see The MDM cluster on page 29.

  Proceed to the section that describes your environment:

  - Replacing a cluster member by adding a new member to the cluster on page 104
  - Replacing a cluster member without adding a new server to the cluster on page 106

Replacing a cluster member by adding a new member to the cluster

This section describes how to replace a member of the cluster, by adding a new member to the cluster to take its place.

Before you begin, perform the following:

- Assign the necessary IP addresses to the replacement server.
- Install the MDM package on the server.

  In this example, we are replacing the server whose IP address is 10.3.1.179, currently a member of a 5-node MDM cluster, with a server a server whose IP address is 10.3.1.57, which is currently external to any VxFlex Ready Node system. This process can be used to replace any role in the MDM cluster.

Procedure

1. Ensure that the current server (179) is not the Master MDM, by running the following command:

   ```
   scli --query_cluster
   ```

   Output, similar to the following, is displayed:

   ```
   # scli --query_cluster
   Cluster:
   Mode: 5_node, State: Normal, Active: 5/5, Replicas: 3/3
   Master MDM:
   Name: mdm17, ID: 0x5d07497754427fd0
   IPs: 10.3.1.17, 192.168.1.17, Management IPs: 10.3.1.17,
   Port: 9011
   Version: 2.0.972
   Slave MDMs:
   ```

```
Name: mdm19, ID: 0x26ee566356362451
IPs: 10.3.1.19, 192.168.1.19, Management IPs: 10.3.1.19,
Port: 9011
Status: Normal, Version: 2.0.972
Name: mdm18, ID: 0x5843c4d16d8f1082
IPs: 10.3.1.18, 192.168.1.18, Management IPs: 10.3.1.18,
Port: 9011
Status: Normal, Version: 2.0.972
Tie-Breakers:
Name: mdm179, ID: 0x7380b70e2f73d346
IPs: 10.3.1.179, 192.168.1.179, Port: 9011
Status: Normal, Version: 2.0.972
Name: mdm20, ID: 0x6dfe1c5f4062b5b3
IPs: 192.168.1.20, 10.3.1.20, Port: 9011
Status: Normal, Version: 2.0.972
```

In this case, server 179 is a Tie Breaker.

2. If the current server is the Master MDM, change its state using the `switch_mdm_ownership` command, as described in the *VxFlex OS CLI Reference Guide*.
3. Make the replacement MDM server a standby MDM, and assign it a name (*mdm57*, in our example) by running the following command, on the Master MDM:

```
scli --add_standby_mdm --mdm_role tb --new_mdm_ip
10.3.1.57,192.168.1.57 --new_mdm_management_ip
10.3.1.57 --new_mdm_name mdm57
```

4. You can see the result of the command by running the following command:

```
scli --query_cluster
```

Output, similar to the following, is displayed:

```
# scli --query_cluster
Cluster:
Mode: 5_node, State: Normal, Active: 5/5, Replicas: 3/3
...
Tie-Breakers:
Name: mdm179, ID: 0x7380b70e2f73d346
IPs: 10.3.1.179, 192.168.1.179, Port: 9011
Status: Normal, Version: 2.0.972
Name: mdm20, ID: 0x6dfe1c5f4062b5b3
IPs: 192.168.1.20, 10.3.1.20, Port: 9011
Status: Normal, Version: 2.0.972
Standby MDMs:
Name: mdm57, ID: 0x073e4c8b1d20d124, Tie Breaker
IPs: 10.3.1.57, 192.168.1.57, Port: 9011
```

mdm57 has been added as a standby MDM. Once it is a standby MDM, it can be added to the cluster.

5. Replace the current *mdm179* with the standby *mdm57* by running the following command:

```
scli --replace_cluster_mdm --remove_tb_name mdm179
--add_tb_name mdm57
```

The following output is displayed:

```
Successfully replaced the cluster MDM
```

The current server has been replaced.

Replacing a cluster member without adding a new server to the cluster

This section describes how to replace a member of the cluster by removing it from the cluster, and then adding it back to the cluster. This procedure requires reducing the amount of nodes in the MDM cluster.

In this example, we are removing the current server whose IP address is 10.3.1.179, currently a Tie Breaker member of a 5-node MDM cluster. Because we must retain a majority in the MDM cluster, we must also remove one of the Slave MDMs in the cluster, in this case the MDM whose IP address is 10.3.1.19. This process can be used to replace any role in the MDM cluster.

Procedure

1. Ensure that the current server (179) is not the Master MDM:

```
scli --query_cluster
```

Output, similar to the following, is displayed:

```
# scli --query_cluster
Cluster:
Mode: 5_node, State: Normal, Active: 5/5, Replicas: 3/3
Master MDM:
Name: mdm17, ID: 0x5d07497754427fd0
IPs: 10.3.1.17, 192.168.1.17, Management IPs: 10.3.1.17,
Port: 9011
Version: 2.0.972
Slave MDMs:
Name: mdm19, ID: 0x26ee566356362451
IPs: 10.3.1.19, 192.168.1.19, Management IPs: 10.3.1.19,
Port: 9011
Status: Normal, Version: 2.0.972
Name: mdm18, ID: 0x5843c4d16d8f1082
IPs: 10.3.1.18, 192.168.1.18, Management IPs: 10.3.1.18,
Port: 9011
Status: Normal, Version: 2.0.972
Tie-Breakers:
Name: mdm179, ID: 0x7380b70e2f73d346
IPs: 10.3.1.179, 192.168.1.179, Port: 9011
Status: Normal, Version: 2.0.972
Name: mdm20, ID: 0x6dfe1c5f4062b5b3
IPs: 192.168.1.20, 10.3.1.20, Port: 9011
Status: Normal, Version: 2.0.972
```

In this case, server 179 is a Tie Breaker.

2. If the current server is the Master MDM, change its state using the `switch_mdm_ownership` command, as described in the *VxFlex OS CLI Reference Guide*.

3. Switch to a 3-node cluster:

```
scli --switch_cluster_mode --cluster_mode 3_node
--remove_tb_name mdm179 --remove_slave_mdm_name mdm19
```

The following output is displayed:

```
Successfully switched the cluster mode.
```

4. To view the result of the command, run:

```
scli --query_cluster
```

Output similar to the following is displayed:

```
# scli --query_cluster
Cluster:
Mode: 3_node, State: Normal, Active: 3/3, Replicas: 2/2
...
Slave MDMs:
Name: mdm18, ID: 0x5843c4d16d8f1082
IPs: 10.3.1.18, 192.168.1.18, Management IPs: 10.3.1.18,
Port: 9011
Status: Normal, Version: 2.0.972
Tie-Breakers:
Name: mdm20, ID: 0x6dfe1c5f4062b5b3
IPs: 192.168.1.20, 10.3.1.20, Port: 9011
Status: Normal, Version: 2.0.972
Standby MDMs:
Name: mdm19, ID: 0x26ee566356362451, Manager
IPs: 10.3.1.19, 192.168.1.19, Management IPs: 10.3.1.19,
Port: 9011
Name: mdm179, ID: 0x7380b70e2f73d346, Tie Breaker
IPs: 10.3.1.179, 192.168.1.179, Port: 9011
```

The cluster has been changed to 3-node mode, as a Slave MDM (*mdm19*) and a TB MDM (*tb179*) have been removed and are now standby MDMs.

Now that the current server is a standby MDM, it can removed from the VxFlex Ready Node system.

5. Remove the current server from the VxFlex Ready Node system:

```
scli --remove_standby_mdm --remove_mdm_name mdm179
```

The following output is displayed:

```
Successfully removed the standby MDM.
```

6. To view the result of the command, run:

```
scli --query_cluster
```

Output similar to the following is displayed:

```
Cluster:
Mode: 3_node, State: Normal, Active: 3/3, Replicas: 2/2
...
Tie-Breakers:
Name: mdm20, ID: 0x6dfe1c5f4062b5b3
IPs: 192.168.1.20, 10.3.1.20, Port: 9011
Status: Normal, Version: 2.0.972
Standby MDMs:
Name: mdm19, ID: 0x26ee566356362451, Manager
IPs: 10.3.1.19, 192.168.1.19, Management IPs: 10.3.1.19,
Port: 9011
```

The current server is no longer a standby MDM.

7. Reassign IP addresses to the current server, as required.

   In our case, we will assign the following IP address to the current server: 10.3.1.57.

8. Add the current server (57) back to the system as a standby MDM, and assign it the name `mdm57`:

```
scli --add_standby_mdm --mdm_role tb --new_mdm_ip
10.3.1.57,192.168.1.57 --new_mdm_management_ip
10.3.1.57 --new_mdm_name mdm57
```

   Output similar to the following is displayed:

```
Successfully added a standby MDM. Object ID 13c925450656db74
```

9. To view the result of the command, run:

```
scli --query_cluster
```

   Output similar to the following is displayed:

```
Cluster:
Mode: 3_node, State: Normal, Active: 3/3, Replicas: 2/2
...
Tie-Breakers:
Name: mdm20, ID: 0x6dfe1c5f4062b5b3
IPs: 192.168.1.20, 10.3.1.20, Port: 9011
Status: Normal, Version: 2.0.972
Standby MDMs:
Name: mdm19, ID: 0x26ee566356362451, Manager
IPs: 10.3.1.19, 192.168.1.19, Management IPs: 10.3.1.19,
Port: 9011
Name: mdm57, ID: 0x13c925450656db74, Tie Breaker
IPs: 10.3.1.57, 192.168.1.57, Port: 9011
```

   The server *mdm57* is now a standby MDM, so it can be promoted to the MDM cluster.

10. Switch to 5-node cluster by adding the standby MDMs to the cluster:

```
scli --switch_cluster_mode --cluster_mode 5_node
--add_slave_mdm_name mdm19 --add_tb_name mdm57
```

The following output is displayed:

```
Successfully switched the cluster mode.
```

11. To view the result of the command, run:

```
scli --query_cluster
```

Output similar to the following is displayed:

```
Cluster:
Mode: 5_node, State: Normal, Active: 5/5, Replicas: 3/3
Master MDM:
Name: mdm17, ID: 0x5d07497754427fd0
IPs: 10.3.1.17, 192.168.1.17, Management IPs: 10.3.1.17,
Port: 9011
Version: 2.0.972
Slave MDMs:
Name: mdm18, ID: 0x5843c4d16d8f1082
IPs: 10.3.1.18, 192.168.1.18, Management IPs: 10.3.1.18,
Port: 9011
Status: Normal, Version: 2.0.972
Name: mdm19, ID: 0x26ee566356362451
IPs: 10.3.1.19, 192.168.1.19, Management IPs: 10.3.1.19,
Port: 9011
Status: Normal, Version: 2.0.972
Tie-Breakers:
Name: mdm20, ID: 0x6dfe1c5f4062b5b3
IPs: 192.168.1.20, 10.3.1.20, Port: 9011
Status: Normal, Version: 2.0.972
Name: mdm57, ID: 0x13c925450656db74
IPs: 10.3.1.57, 192.168.1.57, Port: 9011
Status: Normal, Version: 2.0.972
```

12. When changing an MDM IP address, it is mandatory update and restart all the SDCs in the system as well:

    a. Update the IP addresses:

```
/opt/emc/scaleio/sdc/bin/drv_cfg --mod_mdm_ip
--ip <EXISTING_MDM_IP_ADDRESS>
--new_mdm_ip <NEW_MDM_IP_ADDRESSES>
```

    b. Restart the SDC.

    c. Verify the changes:

```
/opt/emc/scaleio/sdc/bin/drv_cfg --query_mdms
```

Output similar to the following should appear:

```
Retrieved 1 mdm(s)
MDM-ID 043925027bbed30e SDC ID 28c5479b00000000
INSTALLATION ID
7214f7ca647c185b IPs [0]-9.4.4.12 [1]-9.4.4.11
```

Transfer the MDM cluster component using the AMS

Switch an MDM or Tie Breaker (TB) component from one VxFlex Ready Node server to another.

Before you begin

Before moving the cluster component, both the source and the target servers must be part of the VxFlex OS system.

The source is the server from which the MDM or TB will be moved and the target is the server to which the MDM or TB will be moved. If there is no server in the system that can accept the role of an MDM or Tie Breaker component, you must add one before performing this procedure.

Procedure

1. If there is no standby node in the system, add a standby node:
   a. Use the VxFlex OS GUI Hardware view to find a server with a non-MDM role.

      This server will be the target. For example:

Figure 15 Target server

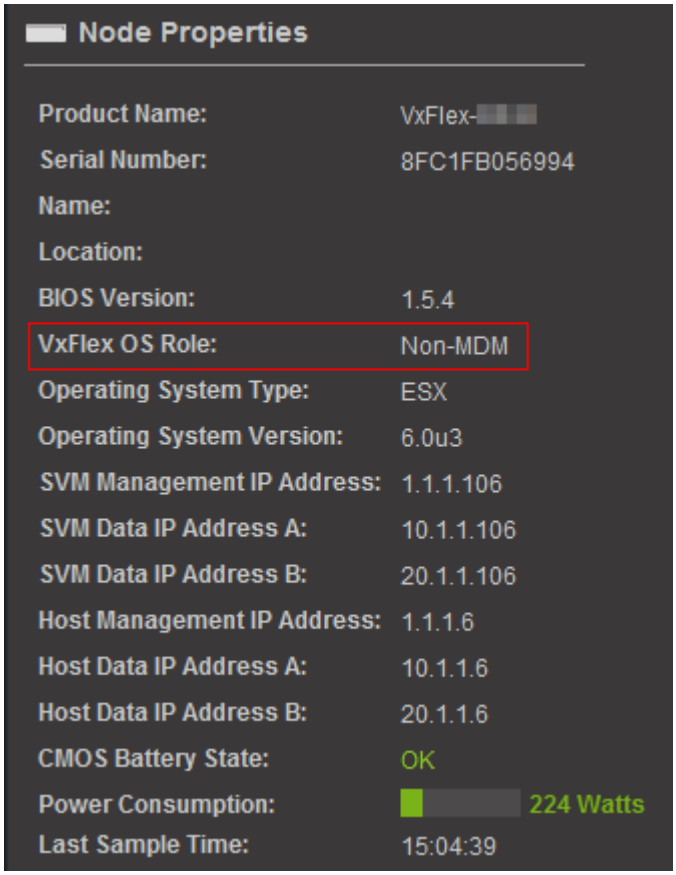

b. In **Hardware** view, right-click the target server and select **Add Standby Node**.

c. In the **Add Standby Node** dialog box, enter a name for the node, select the role that the source server had, and then click **OK**.

2. Right-click the target server and select **Replace Cluster Member**.
3. In the **Replace Cluster Member** dialog box, select the source server (the faulted server) and then click **OK**.
4. When the cluster component has been successfully transferred from the source to the target server, a success message is displayed. Click **Close**.

   The Node Properties sheet displays the source server as a standby MDM and the target server as an MDM in the cluster.

Figure 16 Source server, now a standby MDM

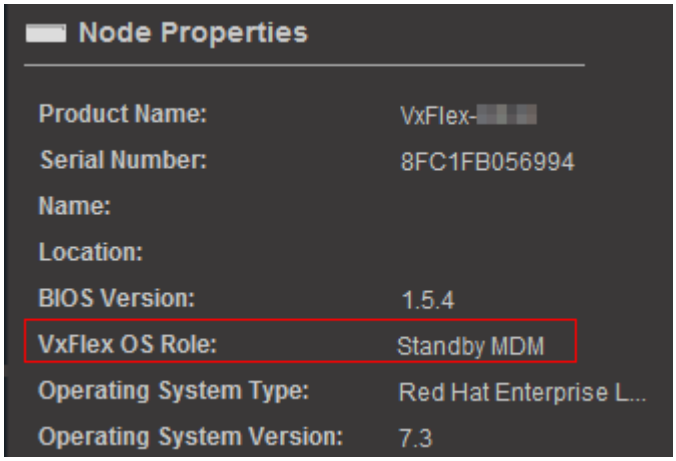

Figure 17 Target server, now an MDM

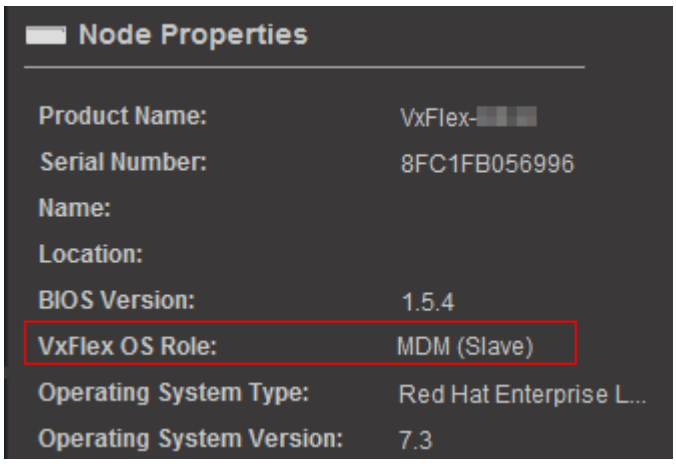

If the process was not successful, contact Customer Support. The virtual IP address on the MDM cluster ensures that all SDCs remain connected to the MDM, even after the physical server is replaced.

Removing a standby MDM with the VxFlex OS GUI

Use the VxFlex OS GUI to remove a Standby MDM.

Before you begin

You must have access to the VxFlex OS GUI.

Procedure

1. From the **Hardware** view, right-click the server and select **Remove Standby**.

Results

This removes the standby from the cluster and removes the MDM.

Configure virtual IP addresses

Configure virtual IP addresses for the MDMs in VxFlex OS.

You can configure virtual IP addresses during deployment or post-deployment. Use the following management tools to configure virtual IP addresses:

Management tool	Actions	Notes
CLI	Add, modify, and remove virtual IP addresses.	For details, see the *VxFlex OS CLI Reference Guide*.
REST API	Add, modify, and remove virtual IP addresses.	For details, see *VxFlex OS REST API Reference Guide*.

Renaming an MDM Cluster Node

Command

rename_mdm

Syntax

```
scli --rename_mdm [--new_mdm_name <NAME> | --i_am_sure]
(--target_mdm_id <ID> | --target_mdm_ip <IP>
[--target_mdm_port <PORT>] | --target_mdm_name <NAME>)
```

Actual command syntax is operating-system dependent. For more information, see "CLI basics".

Description/Notes

Rename an MDM Cluster node. Give MDMs unique names.

Parameters

Parameter	Description
--new_mdm_name <NAME>	The new MDM name (If a new name is not provided, the old name will be deleted.)
--target_mdm_id <ID>	MDM ID
--target_mdm_ip <IP>	MDM IP address
--target_mdm_port <PORT>	MDM port
--target_mdm_name <NAME>	MDM name
--i_am_sure	Preemptive approval

Example

```
scli --rename_mdm --new_mdm_name mdm3 --target_mdm_name
mdm2
```

Managing the SDSs and cache

You can modify and SDS port while there is I/O running. You can also manage the system's cache, including Read Flash cache and Read RAM cache.

Modifying an SDS port during I/O

If you need to modify an SDS port (on-the-fly) while there is I/O running, perform the following steps:

Note

Sometimes, your network topology needs to be prepared for the addition of new port, and this may take some time. VxFlex OS does not prevent unnecessary degraded status or disconnection of SDS. Therefore, in such situations, it is recommended to place the SDS in Maintenance Mode before commencing this procedure.

Procedure

1. On the SDS, perform one of the following, depending on the Operating System:

Operating System	Task
Linux	Run the script: `/opt/emc/scaleio/sds/bin/close_firewall_port.sh`
Windows	From command line, run the batch file `C:\Program Files\EMC\scaleio\sds\bin\close_firewall_port.bat`

2. Open the following SDS file with a text editor, and change the port number shown there to the new port number:

Operating System	File name
Linux	`/opt/emc/scaleio/sds/bin/port`
Windows	`C:\Program Files\EMC\scaleio\sds\bin\port`

3. Open the following SDS configuration file with a text editor:

Operating System	File name
Linux	`/opt/emc/scaleio/sds/cfg/conf.txt`
Windows	`C:\Program Files\EMC\scaleio\sds\cfg\conf.txt`

4. Add the parameter `tgt_port = <NEW_PORT_NUM>` to the file, where <NEW_PORT_NUM> represents the new port number.
5. Perform one of the following:

Operating System	Task
Linux	Run the script: `/opt/emc/scaleio/sds/bin/open_firewall_port.sh`
Windows	From command line, run the batch file `C:\Program Files\EMC\scaleio\sds\bin\open_firewall_port.bat`

6. On the SDS, perform one of the following:

Operating System	Task
Linux	Run the command: `Pkill sds`
Windows	From command line, run the command: `net stop sds_service & net start sds_service`

7. On the MDM, modify the SDS port using the command:

```
scli --modify_sds_port (--sds_id <ID> | --sds_name <NAME> | --
sds_ip <IP>) --new_sds_port <PORT>
```

For example, for an SDS called "sds198" where the new port number is 7071, type:

```
scli --modify_sds_port --sds_name sds198 --new_sds_port 7071
```

Note

If you modify the SDS port on the MDM first, instead of following the above procedure, I/O errors might be encountered.

Managing Read Flash cache

This section describes how to manage the Read Flash Cache (RFcache) feature, which uses PCI flash cards, SSDs and NVMe SSDs for caching of the HDDs in the SDS, thus accelerating the reads of its HDD devices.

RFcache devices and configuration can be performed during initial system deployment. If you want to add the use of the RFcache feature after deployment, use the following work flow:

1. Ensure that the RFcache policy is enabled in the Storage Pool.
2. Enable RFcache in the SDSs where RFcache devices will be added.
3. Add the RFcache devices to the SDSs.

   If you want to stop using the RFcache feature, or remove a specific RFcache device, use the following work flow:
4. Stop RFcache usage in the Storage Pool
5. Remove (command) the RFcache device from the SDS.
6. You can then do one of the following:
   - Physically remove the device from the chassis, and then restart RFcache usage in the Storage Pool

- Add the device to the SDS as a storage device

Managing read RAM cache

This section describes how to manage the read RAM cache feature, which is designed to allocate RAM on nodes for caching of reads or writes. The feature is configured at the following levels:

- For a volume (optional)—specific volume to use or not use caching. If you want all I/Os for a specific volume to use caching, make sure that the volume, the corresponding Storage Pool, and its SDSs are all configured for using caching.
- For one or more Storage Pools—caching must be configured at this level for caching to work in the corresponding SDSs. Write handling mode can also be configured via the `add_storage_pool` command, or in a separate command. To run CLI commands, you must be logged in. Actual command syntax is operating-system dependent. For more information, see the *VxFlex OS CLI Reference Guide*.
- For individual SDSs—caching may be disabled at this level, even if caching in the corresponding Storage Pool is enabled. Cache size can be configured at the SDS level.

  When SDS caching is enabled, the default cache size is 128 MB.

  The amount of RAM you can allocate for cache is limited by the amount of RAM on the SDS server:

- If the RAM is less than 32 GB, 50% of memory can be used for cache
- If the RAM is more than 32 GB, 75% of memory can be used for cache

  The maximum amount of RAM cache is described in Table 2 on page 24.

Note

Only blocks up to 128 k in size will be cached. Any blocks larger than 128k will be ignored by caching.

For a read to be stored in a specific SDS cache, you have to make sure that the cache on that SDS is enabled, and the relevant Storage Pool and the relevant volume are both configured to use cache.

SDC operations

Many SDC operations use drv_cfg. The drv_cfg command line is a local CLI utility that affects only the client on which the SDC is running. Possible SDC operations include updating the SDC driver with IP changes, detecting new volumes, querying volumes, loading a configuration file, adding an MDM, modifying an MDM IP address, enabling support of PDL state, and more.

Updating the SDC driver with IP changes

Procedure

1. Edit `drv_cfg.txt` and change the IP address in the last line to the new IP.

   Location of `drv_cfg.txt`:

   - Linux: `/bin/emc/scaleio/`
   - Windows: In the following registry key - `HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\services\scini\Parameters\mdms`

Enter the IP addresses as a comma-separated list.

Note

On ESXi, GUID and MDM lists are stored as module parameters, and not in a `drv_cfg.txt` file. To modify these parameters, use `esxcli` commands.

2. Save and close the file.
3. Type the following command:

```
/etc/init.d/scini restart
```

Detecting new volumes

Command

`drv_cfg --rescan`

Note

This is not a CLI command, but rather an executable that is run on the SDC server.

Viewing object properties in Backend and Frontend views

Procedure

1. In **Backend** or **Frontend** view, use the filter to display one or more objects, and select the corresponding check boxes of the desired objects (optional).
2. Navigate to the desired object in the table.
3. Display the required information, using the table view options.
4. Select the required object's row in the table, and then, on the expandable Property Sheet on the right side of the window, click the blue arrow buttons beside the headings to expand them and view specific status information.

Note

Contents of the Property Sheet are dynamic, and differ, depending on the row selected in the table. For more information about Property Sheets, see Property Sheets on page 168.

Syntax

```
/opt/emc/scaleio/sdc/bin/drv_cfg --rescan
```

Description/Notes

Volumes are always exposed to the operating system as devices with the prefix `scini` (such as `/dev/scinia`, `/dev/scinib` and so on). Unique names can be found under `/dev/disk/by-id/`.

VxFlex OS periodically scans the system to detect new volumes. You can initiate a scan for the most up-to-date status on a particular SDC node. This command is unique because it is not a CLI command, but rather a command issued on the specific SDC.

Location of `drv_cfg` command:

- Linux: `/opt/emc/scaleio/sdc/bin/drv_cfg`
- Windows: `C:\Program Files\emc\scaleio\sdc\bin\drv_cfg`
- ESXi: Contact Customer Support for access to this tool.

Parameters

Not applicable.

Query volumes using drv_cfg

Command

`drv_cfg --query_vols`

Note

This is not a CLI command, but rather an executable that is run on the SDC server.

Syntax

```
/opt/emc/scaleio/sdc/bin/drv_cfg --query_vols
```

Description/Notes

This utility retrieves information about all known active volume objects in kernel mode. You can use this utility to determine which volumes are mapped, and the ID of each volume in VxFlex OS.

Location of `drv_cfg` command:

- Linux: `/opt/emc/scaleio/sdc/bin/drv_cfg`
- Windows: `C:\Program Files\emc\scaleio\sdc\bin\drv_cfg`
- ESXi: Contact Customer Support for access to this tool.

Example

```
/opt/emc/scaleio/sdc/bin/drv_cfg --query_vols
```

Query tgt objects using drv_cfg

Command

`drv_cfg --query_tgts`

Note

This is not a CLI command, but rather an executable that is run on the SDC server.

Syntax

```
/opt/emc/scaleio/sdc/bin/drv_cfg --query_tgts
```

Description/Notes

This utility retrieves information about all known active tgt objects (SDSs) in kernel mode.

Location of `drv_cfg` command:

- Linux: `/opt/emc/scaleio/sdc/bin/drv_cfg`
- Windows: `C:\Program Files\emc\scaleio\sdc\bin\drv_cfg`
- ESXi: Contact Customer Support for access to this tool.

Example

```
/opt/emc/scaleio/sdc/bin/drv_cfg --query_tgts
```

Query GUID using drv_cfg

Command

`drv_cfg --query_guid`

Note

This is not a CLI command, but rather an executable that is run on the SDC server.

Syntax

```
/opt/emc/scaleio/sdc/bin/drv_cfg --query_guid
```

Description/Notes

This utility retrieves the unique ID of the kernel module. The utility can be used to verify that all SDC GUIDs in the system are unique.

Location of `drv_cfg` command:

- Linux: `/opt/emc/scaleio/sdc/bin/drv_cfg`
- Windows: `C:\Program Files\emc\scaleio\sdc\bin\drv_cfg`
- ESXi: Contact Customer Support for access to this tool.

Note

If the SDC was removed and reinstalled, the GUID of the SDC will be different to its original GUID. In such a case, you may need to remove the SDC, if two SDCs now have the same GUID.

Example

```
/opt/emc/scaleio/sdc/bin/drv_cfg --query_guid
```

Query MDMs using drv_cfg

Command

`drv_cfg --query_mdms`

Note

This is not a CLI command, but rather an executable that is run on the SDC server.

Syntax

```
/opt/emc/scaleio/sdc/bin/drv_cfg --query_mdms
```

Description/Notes

This utility retrieves information about all known MDM objects in kernel mode. This utility is typically used to determine to which MDM an SDC is connected.

Location of `drv_cfg` command:

- Linux: `/opt/emc/scaleio/sdc/bin/drv_cfg`
- Windows: `C:\Program Files\emc\scaleio\sdc\bin\drv_cfg`
- ESXi: Contact Customer Support for access to this tool.

Example

```
/opt/emc/scaleio/sdc/bin/drv_cfg --query_mdms
```

Loading a configuration file using drv_cfg

Command

`drv_cfg --load_cfg_file`

Note

This is not a CLI command, but rather an executable that is run on the SDC server.

This command can not be used on ESX servers. Instead, follow the steps described in "Modifying configuration parameters on ESX servers".

Syntax

```
/opt/emc/scaleio/sdc/bin/drv_cfg
--load_cfg_file <FILE_NAME>
```

Description/Notes

This utility reads a configuration file containing MDM IP addresses, and calls the kernel to connect to them.

Location of `drv_cfg` command:

- Linux: `/opt/emc/scaleio/sdc/bin/drv_cfg`
- Windows: `C:\Program Files\emc\scaleio\sdc\bin\drv_cfg`
- ESXi: Contact Customer Support for access to this tool.

Example

```
/opt/emc/scaleio/sdc/bin/drv_cfg
--load_cfg_file /bin/emc/scaleio/drv_cfg.txt
```

Adding an MDM using drv_cfg

Command

`drv_cfg --add_mdm`

Note

This is not a CLI command, but rather an executable that is run on the SDC server.

This command can not be used on ESX servers. Instead, follow the steps described in "Modifying configuration parameters on ESX servers".

Syntax

```
/opt/emc/scaleio/sdc/bin/drv_cfg --add_mdm
--ip <MDM_IP_ADDRESS_LIST>
```

Description/Notes

This utility calls the kernel module to connect to an MDM. This command is typically used in cases where an SDC is connected to more than one VxFlex OS system.

Location of `drv_cfg` command:

- Linux: `/opt/emc/scaleio/sdc/bin/drv_cfg`
- Windows: `C:\Program Files\emc\scaleio\sdc\bin\drv_cfg`
- ESXi: Contact Customer Support for access to this tool.

Note

Extending your VxFlex OS system with another MDM requires that you update all SDCs in your system with the new MDM IP address. Run the `drv_cfg` utility with the `--mod_mdm_ip` option (see "Modifying an MDM IP address using drv_cfg"), and to make the change persistent, use the `--file` parameter. In addition, any additional objects or systems which interface with the MDM must also be updated.

Parameters

Parameter	Description
--ip <MDM_IP_ADDRESS_LIST>	List of IP addresses (comma delimited) for this Master or Slave MDM
Optional:	
--file <CONFIG_FILE_NAME>	Name of the configuration file to which the MDM information should be written
--only_cfg	Do not call the kernel to actually connect

Example

```
/opt/emc/scaleio/sdc/bin/drv_cfg --add_mdm
--ip 10.100.22.20,10.100.22.30
--file /bin/emc/scaleio/drv_cfg.txt
```

Modifying an MDM IP address using drv_cfg

Command

`drv_cfg --mod_mdm_ip`

Note

This is not a CLI command, but rather an executable that is run on the SDC server. This command can not be used on ESX servers. Instead, follow the steps described in "Modifying configuration parameters on ESX servers".

Syntax

```
/opt/emc/scaleio/sdc/bin/drv_cfg --mod_mdm_ip
--ip <EXISTING_MDM_IP_ADDRESS>
--new_mdm_ip <NEW_MDM_IP_ADDRESSES> [--file <CONFIG_FILE_NAME>]
["-""-"only_cfg]
```

Description/Notes

This utility calls the kernel to modify an MDM's IP address list. It is typically used in cases when an MDM IP address has changed, or when MDMs are added/removed from/to the system. The command must be run on every SDC in the system. To bring the changes into effect, a server restart is required.

Location of `drv_cfg` command:

- Linux: `/opt/emc/scaleio/sdc/bin/drv_cfg`
- Windows: `C:\Program Files\emc\scaleio\sdc\bin\drv_cfg`
- ESXi: Contact Customer Support for access to this tool.

Note

Extending your VxFlex OS system with another MDM requires that you update all SDCs in your system with the new MDM IP address. Ensure that you run this command on all SDCs.

Parameters

Parameter	Description
--ip <EXISTING_MDM_IP_ADDRESS>	One of the existing MDM IP addresses
--new_mdm_ip <NEW_MDM_IP_ADDRESSES>	The new IP address list (comma delimited) for this MDM. If you want to retain the existing address(es), include them in this list.
Optional:	
--file <CONFIG_FILE_NAME>	The name of the configuration file to which the MDM information should be written
--only_cfg	Do not call the kernel to actually connect

Example

```
/opt/emc/scaleio/sdc/bin/drv_cfg --mod_mdm_ip
--ip 10.100.20.20
--new_mdm_ip 10.100.20.20,10.100.20.30,10.100.20.40
```

Permanent Device Loss state

When the MDM has disconnected from the SDC and a volume mapped to this SDC has experienced an I/O error, the ESXi host continuously sends I/Os to the device to determine if the device has become accessible. This can subsequently cause a high I/O-error load that can lead to the host freezing. In cases where the device is disconnected long-term, such as when the entire MDM cluster is down during a network upgrade, the SDC can change the volume state to Permanent Device Loss (PDL) to prevent more I/O errors coming from the ESXi.

PDL is an ESXi state, which, once enabled, is supported in VxFlex OS. Once the ESXi host loses connectivity with a device, if a timeout value is reached, the ESXi will be notified that the device is in a PDL state. The timeout value can be manually set. Once a device is in a PDL state, the ESXi host no longer attempts to re-establish connectivity or issue commands to the device.

Recovering a device from PDL state is described in the VMware documentation for your operating system version. The following link is for ESXi v6.5: https://docs.vmware.com/en/VMware-vSphere/6.5/com.vmware.vsphere.storage.doc/GUID-A513D44C-71DE-47ED-B781-327F78659404.html

Use the instructions that match your environment.

Enabling support of PDL state on the ESXi

Enable support of Permanent Device Loss (PDL) state on the ESXi host.

Procedure

1. On the ESXi host:

```
esxcli system module parameters set -m scini
"<<PREVIOUS_MODULE_PARAMS>> bBlkDevIsPdlActive=1
blkDevPdlTimeoutMillis=<TIMEOUT_VALUE>"
```

   where:

   - *<<PREVIOUS_MODULE_PARAMS>>* is any previous module parameters being used for this ESXi host.
   - *<TIMEOUT_VALUE>* is the timeout time in milliseconds. Its value can be 1000-3600000 (default is 60000) and including it in the command is optional.

2. Reboot the host.

Disabling support of PDL state on the ESXi

Disable support of Permanent Device Loss (PDL) state on the ESXi host.

Procedure

1. On the ESXi host:

```
esxcli system module parameters set -m scini
"<<PREVIOUS_MODULE_PARAMS>> bBlkDevIsPdlActive=0"
```

   where *<<PREVIOUS_MODULE_PARAMS>>* is any previous module parameters being used for this ESXi host.

2. Reboot the host.

Managing ESX servers

Using the following procedures, you can modify parameters on ESX servers and check the SDC state on ESX servers.

Modifying parameters on ESX servers

On an SDC running on an ESX server, esxcli commands can be used in the following cases:

- MDM IP addresses need to be added to the existing list on an SDC
- MDM IP addresses need to be replaced on an SDC
- The SDC's GUID needs to be changed

Specifically, the SDC's GUID or IP address needs to be identified, and then used to add or modify the MDM IP addresses or GUID (depending on the parameter that you want to modify). If you want to add additional MDM IP addresses to existing ones, you must list both old and additional IP addresses in the esxcli command.

For more information about SDC tuning, see *VxFlex OS Performance Fine-Tuning Technical Notes*.

If the current configuration of the VxFlex Ready Node system is registered with a v2.0 VMware plug-in, you can use the plug-in Update SDC parameters to update the MDM IP addresses. For more information, see the *VxFlex Ready Node AMS Deployment Guide*.

Note

These procedures require a server restart to apply the new configuration. The configuration will remain persistent after future server restarts.

To configure MDM IP addresses on the SDC, perform these steps:

Procedure

1. Find the SDC's GUID and the MDM IP addresses configured on the ESX, by typing the command:

```
esxcli system module parameters list -m scini
```

2. In the output of the command, find the existing GUID and MDM IP addresses.

   For example, in the output excerpt below, the GUID and IP addresses are marked in bold:

```
IoctlIniGuidStr string
39b89295-5cfc-4a42-bf89-4cc7e55a1e5b Ini Guid, for example:
12345678-90AB-CDEF-1234-567890ABCDEF
```

```
IoctlMdmIPStr string
9.99.101.22,9.99.101.23 Mdms IPs, IPs for MDM in same cluster
should be comma-separated. To configure more than one cluster
use '+' to separate between IPs.For Example:
10.20.30.40,50.60.70.80+11.22.33.44. Max 1024 characters
```

3. To configure the MDM IP addresses on the SDC, type the command

```
esxcli system module parameters set -m scini -p
"IoctlIniGuidStr=<GUID> IoctlMdmIPStr=<MDM_IPS>"
```

   where *<GUID>* is the existing SDC GUID that you identified in the previous step, and *<MDM_IPS>* is the list of MDM IP addresses. A maximum of 1024 characters is allowed.

   a. To replace the old MDM IP addresses with new MDM IP addresses, omit the old addresses from the command.

   b. To add MDM IP addresses to the existing IP addresses, type both the existing IP addresses and the new IP addresses in the command.

MDM IP addresses for MDMs in same cluster must be comma-separated. To configure more than one cluster, use '+' to separate between IP addresses in different clusters. For example:

```
esxcli system module parameters set -m scini -p
"IoctlIniGuidStr=39b89295-5cfc-4a42-bf89-4cc7e55a1e5b
IoctlMdmIPStr=10.20.30.40,50.60.70.80+11.22.33.44"
```

4. To apply the new configuration, restart the ESX server.

   To change the GUID of the SDC, perform these steps:

5. Find the SDC's GUID and the MDM IP addresses configured on the ESX, by typing the command

```
esxcli system module parameters list -m scini
```

6. In the output of the command, find the existing GUID and MDM IP addresses.

   For example, in the output excerpt below, the GUID and IP addresses are marked in bold:

```
IoctlIniGuidStr string
39b89295-5cfc-4a42-bf89-4cc7e55a1e5b Ini Guid, for example:
12345678-90AB-CDEF-1234-567890ABCDEF
```

```
IoctlMdmIPStr string 9.99.101.22,9.99.101.23 Mdms IPs, IPs for
MDM in same cluster should be comma-separated. To configure
more than one cluster use '+' to separate between IPs.For
Example:
10.20.30.40,50.60.70.80+11.22.33.44. Max 1024 characters
```

7. To change the GUID on the SDC, type the command

```
esxcli system module parameters set -m scini -p
"IoctlIniGuidStr=<NEW_GUID> IoctlMdmIPStr=<MDM_IPS>
```

   where <NEW_GUID> is the new SDC GUID, and <MDM_IPS> is the list of MDM IP addresses that you identified in the previous step. You must include these IP addresses in the command.

   For example:

```
esxcli system module parameters set -m scini -p
"IoctlIniGuidStr=28a78184-4beb-4a42-bf89-4cc7e55a1e5b
IoctlMdmIPStr= 9.99.101.22,9.99.101.23"
```

8. To apply the new configuration, restart the ESX server.

Checking the SDC state on ESX servers

On an SDC running on an ESX server, an esxcli command can be used to check the current state of the SDC.

To display the SDC state on the ESX server, type the following command:

```
esxcli system module list |grep scini
```

The following examples show typical outputs of the command:

- Output where driver is installed and enabled to load, but not loaded:

```
Name Is Loaded Is Enabled
---------------------------- --------- ----------
scini false true
```

- Example of SDC in correct state (enabled and loaded):

```
Name Is Loaded Is Enabled

---------------------------- --------- ----------
scini true true
```

CHAPTER 6

Security Management

The following topics describe security management in VxFlex Ready Node.

- Sign and update AMS security certificates 130
- Signing AMS security certificates using CLI commands 131
- Modifying AMS-MDM connection security 132
- Modifying AMS-MDM connection security 134
- Back up SSH public keys using the VxFlex OS GUI 136
- Backing up SSH public keys using CLI commands 137
- Changing the emergency password 138
- Setting credentials for existing Monitor users 138
- Changing login credentials for existing Management users 139
- Login banner overview 140

Sign and update AMS security certificates

A security certificate can be signed by exporting it from VxFlex OS, getting it signed by the Certificate Authority (CA) server, and then importing it back into the system. Afterwards, the **Renew Certificates** option updates the certificates on the VxFlex OS components.

Before you begin

Ensure that you can authenticate to the VxFlex OS GUI with an admin user.

The following procedure explains how to export a certificate from VxFlex OS, import it back into the system after it has been signed, and then how to update VxFlex OS components with the signed certificate.

Note

This procedure is optional, and can be done before or after deployment.

Procedure

1. From the **System Settings** menu at the top right side of the window, in any view, select the **System Settings** option.

   The **System Settings** window is displayed.
2. In the left panel, click **Connection**.
3. Export a certificate:
   a. Scroll down to the **Export AMS Certificate** section, set the required **Certificate Key Length**, and click **Get Certificate**.

      A certificate file is created in CSR format, and a **Save** window is displayed.

   b. Navigate to a location of your choice, and save the file.
4. On your CA server, sign the certificate (CSR) file.
5. Import the signed CSR file:
   a. In the VxFlex OS GUI, return to the **System Settings** window, **Connection** pane.

   b. Scroll down to the **Export AMS Certificate** section, and click **Import Certificate**.

      An **Open** window is displayed.

   c. Navigate to the signed CSR file, and click **Open**.

      The signed certificate is imported into the system.
6. In the VxFlex OS GUI, from the **System Settings** menu at the top right side of the window, select the **Renew Certificates** option.

   The certificate renewal process will begin updating the certificates in VxFlex OS components.

Signing AMS security certificates using CLI commands

Sign the AMS security certificates that are generated by VxFlex OS components, using CLI commands.

The Certificate Authority (CA) server is the signing authority for signing and validating the security certificates. For the CA server to sign a security certificate, you must complete the following sequential stages:

1. Exporting the certificate from the VxFlex OS system
2. Getting the certificate signed by your own CA server
3. Importing the certificate back into the system
4. Updating the certificate on VxFlex OS components, via Renew Certificates options

Note

This procedure is optional and can be performed either before VxFlex OS deployment or post deployment.

Procedure

1. Log in to the AMS.
2. Generate a certificate from the Linux-based AMS server:
   a. Perform either of the following:
      - Access full path: `java -jar /opt/emc/scaleio/ams/webapps/ROOT/resources/scaleio_repository/`
      - Relative path: `cd /opt/emc/scaleio/ams/webapps/ROOT/resources/scaleio_repository`

   b. Run this command:

   ```
   EMC-ScaleIO-ams-cli-2.6-<build>.XXX.jar --username
   <USERNAME> --password <PASSWORD>  --
   generate_certificate_request --file_name <NAME_FILE>
   ```

   An unsigned certificate appears in the command-defined file on your local drive.
3. Send the command-defined file, with unsigned certificate, to your CA server and initiate request for the signed certificate.
4. Upload the signed certificate to the server.
5. Replace the existing unsigned certificate in the Linux-based AMS server with the new signed certificate:
   a. Perform either of the following:
      - Access full path: `java -jar /opt/emc/scaleio/ams/webapps/ROOT/resources/scaleio_repository/`
      - Relative path: `cd /opt/emc/scaleio/ams/webapps/ROOT/resources/scaleio_repository`

b. Run this command:

```
EMC-ScaleIO-ams-cli-2.6-<build>.XXX.jar --username
<USERNAME> --password <PASSWORD>  --
replace_ams_certificate  --file_name <NAME_FILE>
```

The signed certificate replaces the unsigned certificate in the AMS server.

6. Start the process of renewing the certificate in the Linux-based AMS server:

   a. Perform either of the following:

   - Access full path: `java -jar /opt/emc/scaleio/ams/webapps/ROOT/resources/scaleio_repository/`
   - Relative path: `cd /opt/emc/scaleio/ams/webapps/ROOT/resources/scaleio_repository`

   b. Run this command:

```
EMC-ScaleIO-ams-cli-2.6-<build>.jar --username <USERNAME> --
password <PASSWORD>  --replace_nodes_certificates
```

The signed certificate replaces the unsigned certificate in the AMS server.

Modifying AMS-MDM connection security

Set up secure or non-secure connectivity between AMS and MDM, using the VxFlex OS GUI. The default setting is secure AMS-MDM connectivity.

Before you begin

Ensure that you have:

- The VxFlex OS GUI installed on your computer
- The username and password for logging in to the VxFlex OS GUI

It is recommended to always use a secure AMS-MDM connection for accessing the VxFlex OS MDM. Only when access to the MDM is denied, for reasons such as expiration of a security certificate, and you need MDM access for maintenance, should you disable the secure AMS-MDM connection. You should fix the issue, and then enable the secure AMS-MDM connection again.

Procedure

1. Log in to the VxFlex OS GUI.
2. Display the **System Settings** menu.

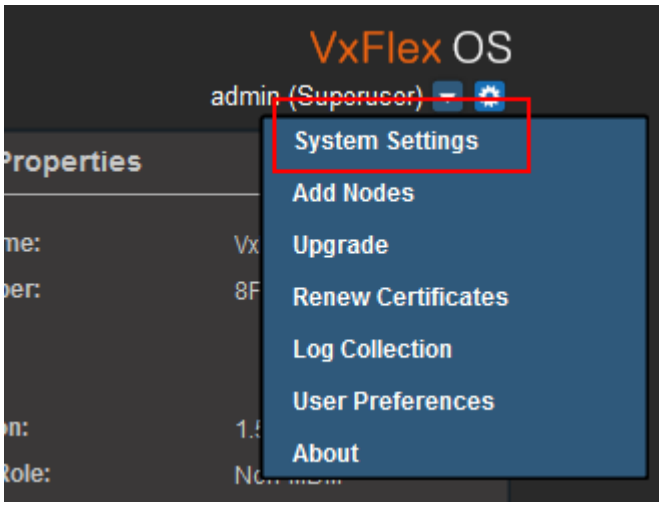

3. Select **System Settings**.
4. In the **System Settings** navigation pane, select **Connection**.
5. Scroll down to the **AMS Secured Connectivity Actions** section.

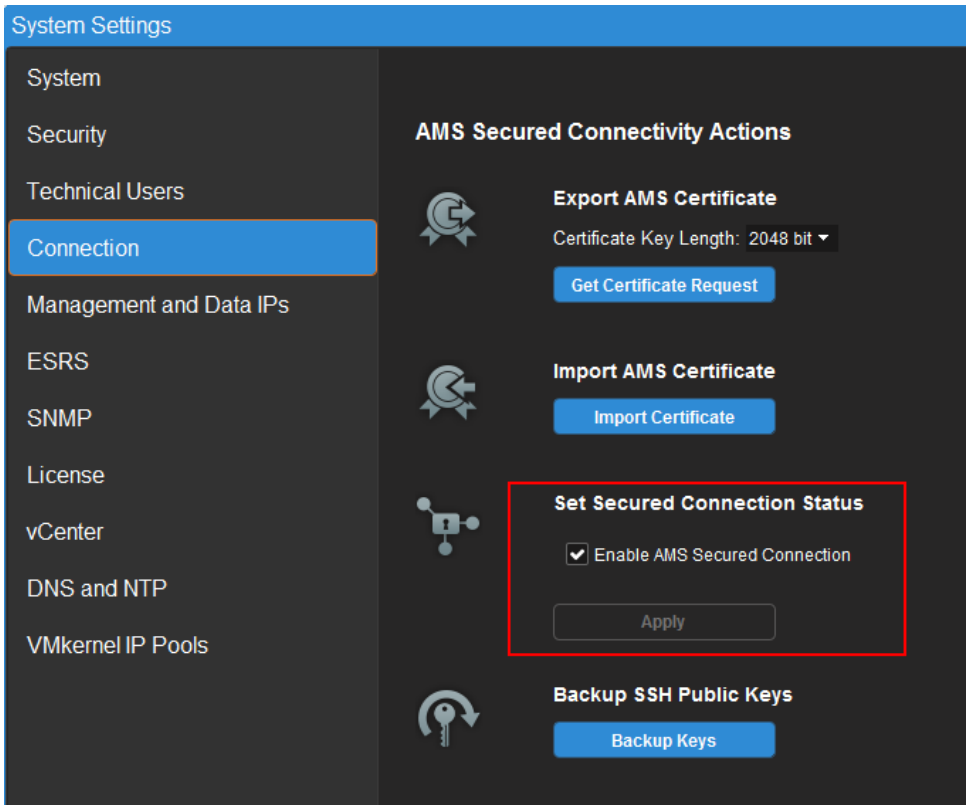

6. Under **Set Secured Connection Status** option, do any of the following (depends on the default secure/non-secure setting):
   - To disable AMS-MDM secure connectivity (to move from secure to non-secure connectivity):
     a. Verify that the **Enable AMS Secured Connection** checkbox is selected (default selection).
     b. Deselect the **Enable AMS Secured Connection** checkbox.

c. Click **Apply**.

d. Click **Close** to close the screen.

e. In the VxFlex OS GUI, select the **Monitor > Alerts** tab to display the list of alerts.

f. Verify that the alert `Secure Connection disabled on MDM` is triggered.

- To enable AMS-MDM secure connectivity (to move from non-secure to secure connectivity):

  a. Verify that the **Enable AMS Secured Connection** checkbox is deselected.

  b. Select the **Enable AMS Secured Connection** checkbox.

  c. Click **Apply**.

  d. Click **Close** to close the screen.

  e. On the VxFlex OS GUI, select the **Monitor > Alerts** tab to display the list of alerts.

  f. Verify that the alert `Secure Connection disabled on MDM` is removed.

Modifying AMS-MDM connection security

Set up the secure or non-secure connectivity between AMS and MDM, using the AMS CLI. The default setting is secure AMS-MDM connectivity.

Before you begin

Ensure that you have:

- The user name and password for logging in to the AMS server

It is recommended to always use a secure AMS-MDM connection for accessing the VxFlex OS MDM. Only when access to the MDM is denied, for reasons such as expiration of a security certificate, and you need MDM access for maintenance, should you disable the secure AMS-MDM connection. You should fix the issue, and then enable the secure AMS-MDM connection again.

Procedure

1. Log in to the AMS server.
2. Perform one of the following:
   - To change connectivity from secure mode to non-secure mode, perform the following steps:

Option	Description
Windows	a. Open a command line session. b. • To run the full path command, type: `java -jar C:\Program Files\EMC\scaleio\AMS\webapps \ROOT\resources\scaleio_repository\EMC-ScaleIO-ams-`

Option	Description

	`cli-<version>.jar --set_ams_secure_mode disable_secure_mode`
	• To run the command from a relative path, type:
	`cd "C:\Program Files\EMC\scaleio\AMS\webapps\ROOT\resources\scaleio_repository"`
	and then:
	`java -jar EMC-ScaleIO-ams-cli-<version>.jar --set_ams_secure_mode disable_secure_mode`

Linux	a. Open a command line session.
	b. • To run the full path command, type:
	`java -jar /opt/emc/scaleio/ams/webapps/ROOT/resources/scaleio_repository /EMC-ScaleIO-ams-cli-<version>.jar --set_ams_secure_mode disable_secure_mode`
	• To run the command from a relative path, type:
	`cd /opt/emc/scaleio/ams/webapps/ROOT/resources/scaleio_repository`
	and then:
	`java -jar EMC-ScaleIO-ams-cli-<version>.jar --set_ams_secure_mode disable_secure_mode`

- To change connectivity from non-secure mode to secure mode, perform the following steps:

Option	Description

Windows	a. Open a command line session.
	b. • To run the full path command, type:
	`java -jar C:\Program Files\EMC\scaleio\AMS\webapps\ROOT\resources\scaleio_repository\EMC-ScaleIO-ams-cli-<version>.jar --set_ams_secure_mode enable_secure_mode`
	• To run the command from a relative path, type:
	`cd "C:\Program Files\EMC\scaleio\AMS\webapps\ROOT\resources\scaleio_repository"`

Option	Description
	and then: `java -jar EMC-ScaleIO-ams-cli-<version>.jar --set_ams_secure_mode enable_secure_mode`

Linux	a. Open a command line session.
	b. • To run the full path command, type:
	`java -jar /opt/emc/scaleio/ams/webapps/ROOT/resources/scaleio_repository /EMC-ScaleIO-ams-cli-<version>.jar --set_ams_secure_mode disable_secure_mode`
	• To run the command from a relative path, type:
	`cd /opt/emc/scaleio/ams/webapps/ROOT/resources/scaleio_repository`
	and then:
	`java -jar EMC-ScaleIO-ams-cli-<version>.jar --set_ams_secure_mode enable_secure_mode`

Back up SSH public keys using the VxFlex OS GUI

The VxFlex OS GUI can be used to back up SSH public keys.

Before you begin

Ensure that you can authenticate to the VxFlex OS GUI with an admin user.

It is recommended that you back up SSH public keys to your local drive for recovery purposes.
When using secure mode during VxFlex OS system recovery process, it is recommended that you back up SSH public keys after you perform system deployment or add/remove a server node. If you do not backup SSH public keys, the recovery process can run in non-secure mode as well.

However, for hardened system security, it is recommended that you always use SSH public keys for system access.

Procedure

1. From the **System Settings** menu at the top right of the window, in any view, select the **System Settings** option.

   The **System Settings** window is displayed.
2. In the navigation pane, click **Connection**.
3. Scroll down to the **Backup SSH Public Keys** section, and click **Backup Keys**.
4. Accept the default location and file name, or modify them as required. Make a note of these for future reference.

   The default file name is `<AMS_IP_address>_ssh_keys.txt`.

5. Click **Save**.

Results

When the process is finished, the VxFlex OS GUI displays a confirmation message.

Backing up SSH public keys using CLI commands

Back up SSH public keys in a VxFlex OS AMS using CLI commands.

Before you begin

Ensure that you have user name and password for accessing the AMS server.

It is recommended that you back up SSH public keys to your local drive, for recovery purposes.
When using secure mode during a VxFlex OS system recovery process, it is recommended that you back up SSH public keys after you perform system deployment or add/remove a server node.

If you do not backup SSH public keys, you can run the Recovery process in non-secure mode. However, for the reason of system security, it is recommended that you always use SSH public keys for system access.

Procedure

1. Log in to the AMS server.
2. For Linux as the AMS operating system, do any of the following:
   - Backing up SSH keys via file access:
     a. Run the following command:

```
java -jar /opt/emc/scaleio/ams/webapps/ROOT/resources/
scaleio_repository EMC-ScaleIO-ams-cli-XXX.jar --
username <Username>  --password <Password> --
export_trusted_ssh_keys --file_name WANTED_FILE_NAME
```

     Output, similar to the following, appears:

```
{

  "status" : "Accepted"

}
```

   - Backing up SSH keys via folder access:
     a. Go to the folder: `Cd /opt/emc/scaleio/ams/webapps/ROOT/resources/scaleio_repository`
     b. Run the fallowing command:

```
java -jar EMC-ScaleIO-ams-cli-XXX.jar --username
<Username>  --password <Password> --
export_trusted_ssh_keys --file_name WANTED_FILE_NAME
```

     Output, similar to the following, appears:

```
{

  "status" : "Accepted"
```

```
}
```

Results

The backup of SSH public keys for all system nodes appears on your local drive, in the file set up in the command (in this example, WANTED_FILE_NAME).

Changing the emergency password

Using the VxFlex OS GUI, change the emergency password used for offline access to the VxFlex OS GUI.

Before you begin

To perform this procedure, you must know the existing emergency password.

The emergency password is defined during system deployment. It can be used to log in to the VxFlex OS GUI when there is no connectivity to the system's MDM. The following procedure explains how to change the emergency password.

Procedure

1. From the **System Settings** menu at the top right side of the window, select the **System Settings** option.

   The **System Settings** window is displayed.
2. In the left panel, click **Technical Users**.

   The **Set Credentials for Existing Technical Users** panel is displayed on the right side of the window.
3. In the **Current Password** box, type the password currently defined as the emergency password.
4. In the **New Password** and **Confirm Password** boxes, type the new emergency password.
5. Click **Save**.

Results

The emergency password has been changed.

Setting credentials for existing Monitor users

Using the VxFlex OS GUI, configure credentials for a single local Monitor user, or for an LDAP group with Monitor privileges.

Before you begin

Administrator privileges are required to configure the local Monitor user.

If the Monitor user is an LDAP group, ensure that the LDAP group is also configured with Monitor privileges.

Ensure that you can authenticate to the VxFlex OS GUI with an admin user.

Configure Monitor User settings to create login credentials for either a single local user with Monitor-only access privileges, or for an LDAP group with Monitor access privileges. The default setting is the Admin user.

Procedure

1. From the **System Settings** menu at the top right side of the window, select the **System Settings** option.

   The **System Settings** window is displayed.
2. In the left panel, click **Technical Users**.

   The **Set Credentials for Existing Technical Users** panel is displayed on the right side of the window.
   Figure 18 System Settings window, Technical Users pane, Monitor User
3. In the **User name** box, type the local or LDAP user name that has Monitor privileges.
4. In the **Password** box, type the password.
5. Click **Save**.

Results

The Monitor User credentials have been defined in the system.

Changing login credentials for existing Management users

Use the VxFlex OS GUI to change login credentials for the host Operating System user, SVM user, and BMC (iDRAC) user.

Before you begin

The host operating system user, SVM user, and BMC (iDRAC) user credentials are configured during initial system deployment. If you need to change any of them, do so using the VxFlex OS GUI.

Procedure

1. From the **System Settings** menu at the top right side of the window, select the **System Settings** option.

   The **System Settings** window is displayed.
2. In the left panel, click **Technical Users**.

   The **Set Credentials for Existing Technical Users** panel is displayed on the right side of the window.
3. Click **Management Users**.

   The credentials currently defined for management users are displayed.
4. For the required management user type, type the new username and password in the corresponding boxes.
5. Click **Save**.

Results

The changes that you made to the credentials are saved in the system.

Login banner overview

A login banner is a text file that is displayed upon login to the system. It can be used to communicate messages or to obtain user consent to real-time monitoring of information and retrieval of stored files.

When the login banner is set up, it appears during the system login process before the login credential prompts. The login banner displays differently in the VxFlex OS GUI and CLI interfaces:

- VxFlex OS GUI:
  - When logging in, the login banner is displayed, and must be approved.
- CLI:
  - When logging in, the user is prompted to press any key, after which the banner is displayed.
  - To continue, the banner must be approved.

Limitations:

- Only users with administrative security rights can set up, update, or remove the login banner.
- Text files up to 16 bytes are supported.
- Only one login banner is supported.

Setting up a login banner using the CLI

Use the CLI to set up, modify, or stop displaying a login banner.

Before you begin

Ensure that you have access to the IP address of the Master MDM.

Procedure

1. Log in to VxFlex OS using the IP address of the Master MDM.
2. Perform the desired operation:

Option	Description
Create (or modify) a new banner	a. Create a text file (or modify an existing file) with the message that you want to display in the login banner. b. Run the following command: `scli --set_login_banner --filename <FILENAME>` where *<FILENAME>* is the path of the login banner text file. The login banner is displayed the next time a user logs in to VxFlex OS.

Option	Description
Stop displaying the banner	a. Run the following command: `scli --set_login_banner --remove_banner`

Enabling or disabling preemptive acceptance of the login banner

Preemptive acceptance of the login banner allows the user to bypass the login banner, for example, when running scripts. A user with admin security rights can enable or disable the option of preemptive acceptance. By default, preemptive acceptance is enabled and the login banner can be bypassed using a CLI command.

Before you begin

To enable or disable the preemptive acceptance option, you must have administrative rights.

Procedure

1. Log in to VxFlex OS:

```
scli --login --username admin --password <PASSWORD>
```

2. Run the following command to enable preemptive acceptance:

```
scli --set_cli_login_banner_preemptive_acceptance --enable
```

3. Run the following command to disable preemptive acceptance:

```
scli --set_cli_login_banner_preemptive_acceptance --disable
```

Activating preemptive acceptance of the login banner

When preemptive acceptance of the login banner is enabled (default), you can log in to VxFlex OS in a special way that activates preemptive acceptance of the login banner.

Before you begin

Preemptive acceptance of the login banner is enabled.

Procedure

1. Log in to VxFlex OS with the `accept_banner_by_scripts_only` parameter:

```
scli --login --username <USERNAME> --
accept_banner_by_scripts_only
```

where *<USERNAME>* is the user running the script.

CHAPTER 7

Opening the VxFlex OS GUI and Logging In

The following topics describe how to open the VxFlex OS GUI, and log in procedures

- Log in to the VxFlex OS GUI 144
- Connection and disconnection information 144

Log in to the VxFlex OS GUI

Open and log in to the VxFlex OS GUI.

Before you begin

Ensure that:

- The VxFlex OS GUI software is installed on the workstation. To install the VxFlex OS GUI, see "Install the VxFlex OS GUI."
- You have these credentials (available from the administrator):
  - AMS management IP address or hostname
  - Username (default: admin)
  - Password (defined during deployment)

Procedure

1. Open the VxFlex OS GUI by selecting **Start** > **All Programs** > **EMC** > **VxFlex OS GUI**.

   The initial login screen is displayed.
2. Type the IP address or hostname and click **Connect**.

   If a certificate notice is displayed, review and accept the certificate.

   If a login banner is displayed, confirm it to continue.
3. In the login screen, type the username and password, and click **Login**.

Results

The VxFlex OS GUI is displayed.

After you finish

Users and passwords are configured with the VxFlex OS CLI. For more information, see the "Security" chapter of the *VxFlex OS Ready Node User Guide*.

Connection and disconnection information

You can check at any time to which IP address your VxFlex OS GUI is connected, using the following methods:

- View the IP address displayed in the top left corner of the VxFlex OS GUI window.
- Hover your mouse pointer over the **Management** tile on the Dashboard. A tooltip displays connection information for the nodes in the MDM cluster, and the management IP addresses
  If the VxFlex OS GUI loses its connection with the MDM, the window display is dimmed, and a notification dialog box is displayed.

CHAPTER 8

VxFlex OS GUI Features

The following topics describe VxFlex OS GUI features.

- VxFlex OS GUI overview 146
- VxFlex OS GUI conventions 147
- Hardware view 148
- Dashboard view 150
- Frontend views 157
- Backend view 161
- Alerts view 167
- Property Sheets 168

VxFlex OS GUI overview

Use the VxFlex OS GUI to monitor and configure VxFlex OS.

The various windows display different views and data that are beneficial to the storage administrator. You can review the overall status of the system, drill down to the object level, and monitor these objects. You can use the VxFlex OS GUI to provision and modify many of the objects.

The following sections in this chapter describe the available windows, and how to use them. The Glossary at the end of this publication provides more detailed information about the objects and properties displayed in the VxFlex OS GUI. Your user privileges control the features displayed by the VxFlex OS GUI. If you cannot access certain commands, check whether you have the appropriate user permissions, and if necessary, contact your system administrator for assistance.

The following table provides a general overview of the tasks that you can perform with the VxFlex OS GUI.

Table 24 VxFlex OS GUI task overview

To do this...	Use...
See a general overview of the entire system	Hardware view—see Hardware view on page 148
See detailed information about one or more backend system objects, in table format (filtering available)	Backend view—see Backend view on page 161
See detailed information about one or more frontend system objects, in table format (filtering available)	Frontend views—seeFrontend views on page 157
See very detailed information about a specific system object	Property Sheets—see Property Sheets on page 168
Minimize the main window to a floating widget	Widget—see Widget on page 152
See a list of errors and alerts currently active in the system	Alerts view—see Alerts view on page 167
Add, remove, maintain, or configure a system object (backend)	Backend view—see Configuring the System using the VxFlex OS GUI on page 179
Add, remove, or configure a system object (hardware, caching, SDS, devices)	Hardware view—seeHardware view on page 148, and Configuring the System using the VxFlex OS GUI on page 179
Add, remove, or configure volumes, snapshots or SDCs	Frontend views—see Configuring the System using the VxFlex OS GUI on page 179
Monitor various aspects of your system	Backend and Frontend views—see Monitoring the System using the VxFlex OS GUI on page 173 Hardware view—see Viewing detailed hardware properties on page 174
Find information about your system's license	System Settings menu—see Viewing licensing information on page 177

Table 24 VxFlex OS GUI task overview (continued)

To do this...	Use...
Modify system settings, such as ESRS, SNMP, passwords and license configurations; discovering new system components	System Settings menu— see Modifying system settings on page 238

VxFlex OS GUI conventions

This section describes conventions used in the VxFlex OS GUI, including alert indicators and color codes.

Alerts indicators

The Alerts indicators show the overall error state of the system. When lit, indicators show the number of active alerts of each severity. Similar indicators are displayed in some views of the Backend table, and also on Property Sheets (in some cases, an additional blue indicator for information only is included). You can view details about the alerts active in the system in the Alerts view. For more information about Alerts, see Alerts view on page 167 and VxFlex OS Alerts in SNMP, VxFlex OS GUI, REST, and ESRS on page 283.

Figure 19 Alerts indicators

Color codes

Color codes provide quick visual feedback on the status of various objects in the system. The following tables summarize the colors used in the system, and their meaning. The color codes are used in a variety of elements and views in the user interface.

Table 25 VxFlex OS GUIcolor codes

Color	Meaning	Dashboard View	Backend View
LIGHT GREEN (protected)	Available protected, healthy storage	✓	✓
GREEN (in maintenance)	One or more SDSs are in Maintenance Mode, and part of those SDSs' capacity is temporarily protected on other SDSs	✓	✓
YELLOW	Snapshot capacity (yellow outline)	✓	
	Capacity-related statuses		✓
ORANGE (degraded)	Data is not protected	✓	✓
	Rebuild or Rebalance in progress	✓	

Table 25 VxFlex OS GUIcolor codes (continued)

Color	Meaning	Dashboard View	Backend View
RED (failed)	Data is unavailable	✓	✓
DARK GRAY	Unused capacity No activity or zero values	✓	 ✓
DARK GRAY striped with RED	System is unable to determine if capacity is Unavailable or Unused	✓	✓
PALE GRAY	Decreased capacity. This capacity exists physically, but has been disabled (typically to allow maintenance tasks on devices).	✓	
BLUE (spare)	Capacity reserved for recovery purposes	✓	
BRONZE	Volume capacity	✓	
DARK BLUE	Indicates selected items in the filter		✓

Table 26 Alert symbols and color codes

Color	Meaning	Dashboard View	Backend View	Alert View
YELLOW	Low alert status	✓	✓	✓
ORANGE	Medium alert status	✓	✓	✓
RED	High alert status	✓	✓	✓
LIGHT BLUE	Information message (no faults)		✓	✓

Hardware view

The **Hardware** view shows the status of the physical hardware components in your node. This view provides read-only information, and lets you perform some configuration activities, such as adding and removing physical storage devices, adding and removing devices used for caching purposes, and more.

Figure 20 Hardware view

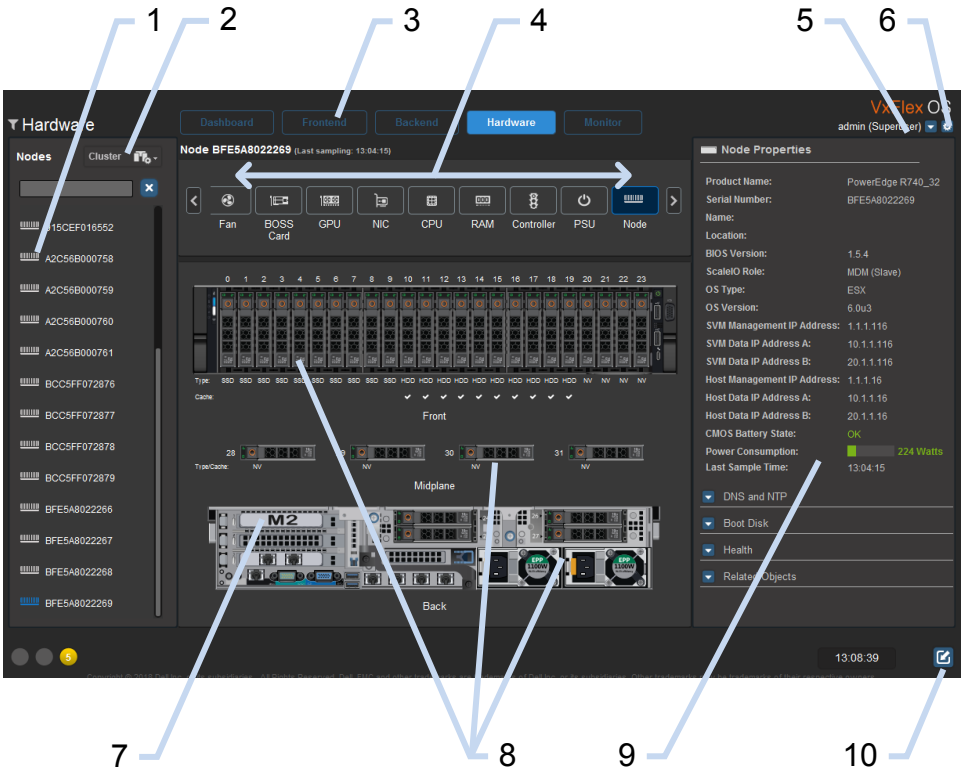

1 **Nodes**—Displays a list of all the nodes in the system. Information about the selected node is displayed in the middle and right panes of the window.

2 **Cluster** menu—contains commands used to perform a cluster operation on a selected node in the nodes list.

3 **Tabs**—Display various views in the VxFlex OS GUI; some tabs provide further options when clicked

6 **System Settings** menu—System Settings, Add Nodes, Upgrade, Log Collection, Renew Certificates, User Preferences, About

7 **BOSS card**—*M2* shown on the back panel represents the BOSS card.

8 **Detailed components**—Display details about the selected component in the property sheet on the right. Error states are indicated by a colored border around the relevant component. Devices have context menus containing commands. Below devices, device type and caching configuration are shown. Check marks below devices indicate that the devices are being cached.

Note

Some models contain devices in the chassis midplane. In these cases, the midplane devices are shown in the VxFlex OS GUI in the area between the front and back panels.

4 **Hardware components**—Each button displays relevant information in the **Property Sheet** in the panel on the right side of the window. Error states are indicated by a colored border around the relevant button.

5 **User Name**—name of user currently logged in to the VxFlex OS GUI. To the right of the user name is the **Logout** menu.

9 **Property Sheets**—Provide detailed read-only information about the object selected in the middle panels of the window

10 **Widget toggle**—Minimizes the display to a widget version of the Dashboard

Dashboard view

The Dashboard displays the overall system status. Each tile displays a certain aspect of the storage system. Various controls let you customize the information displayed on the Dashboard. The following figure shows the Dashboard controls.

Figure 21 Dashboard controls

1 **Navigation toggle**—Displays navigation tree

2 **Current view**—Displays object currently being applied to the dashboard

6 **System Settings** menu—System Settings, Add Nodes, Upgrade, Log Collection, Renew Certificates, User Preferences, About

7 **Color code legend**—Displays the legend

3	**Dashboard**—Displays the Dashboard view, currently selected	8	**I/O Workload toggles**—control the information displayed in this tile. For more information, see Customizing system preferences on page 235.
4	Other tabs—Display various views in the VxFlex OS GUI; some tabs provide further options when clicked	9	**Dimmer toggle**—Enables\disables the dimmer
5	**Username**—name of user currently logged in to the VxFlex OS GUI. To the right of the username is the **Logout** menu.	10	**Widget toggle**—Minimizes the display to a widget

Navigation tree

The Dashboard's navigation button toggles the display of the navigation tree. The navigation tree is hierarchical, and controls the Dashboard display. You can display information on the Dashboard according to:

- Entire system (default)
- Protection Domain
- Storage Pool

You can change the Dashboard display by double-clicking the desired navigation tree node. Some tiles on the Dashboard may be dimmed if they are not relevant for the node that you have selected.

System Settings and Logout menus

The **Logout** menu lets you do the following:

- View the user name of the user currently logged in to the system (default)
- Log out of the system

The **System Settings** menu lets you do the following:

- Open the **System Settings** window, where you can configure and modify security, IP addresses, ESRS and SNMP connectivity, user passwords, licensing, vCenter and discovery of new components. For more information, see Modifying system settings on page 238.
- Add nodes to the system, as described in the *VxFlex Ready Node AMS Deployment Guide.*
- Upgrade system components, as described in the *VxFlex Ready Node AMS Upgrade Guide*.
  The non-disruptive operation can upgrade the following components:
  - VxFlex OS components (MDM and TieBreaker, SDS, SDC, LIA, and RFcache)
  - DAS Cache
  - Operating system patches
  - VxFlex OS AMS and VxFlex OS GUI
  - BMC (iDRAC)
  - Node firmware (Storage controller, NIC, disk)

- BIOS

All VxFlex OS software and hypervisor upgrades (including security patches) in the AMS environment must be orchestrated by AMS, as described in the *VxFlex Ready Node AMS Upgrade Guide*. Upgrades performed using any method other than AMS may impact the AMS management capabilities in general, and orchestrated upgrades in particular.

Dimmer

The **Dimmer** button toggles the dimmer feature on and off. When you use the dimmer, only tiles that are essential for real-time monitoring are lit. You can temporarily light up the dimmed tiles by hovering the mouse pointer over them.

Widget

The **Widget** button reduces the dashboard size to a widget containing a condensed display of the storage system. The widget floats on your desktop, allowing you to visually monitor your system while using other applications. The widget displays the **Capacity** tile, Workload activity, Rebuild/Rebalance activity, and Alerts indicators.

When the dashboard is minimized to a widget, the **Full-Screen** button is displayed in the top-right-corner, as shown in the following figure. This button toggles the display back to full dashboard mode.

Figure 22 Widget, showing Return to Full-Screen button

Dashboard tiles

The Dashboard tiles provide a visual overview of storage system status. The tiles are dynamic, and contents are refreshed at the interval set in the system preferences (default: 10 second intervals). System preferences can also be used to set the display to basic or advanced reporting. For more information, see Customizing system preferences on page 235.

Some of the tiles' contents differ, depending on the navigation filter in use. When the dimmer feature is enabled, the display of non-essential tiles is dimmed, unless the mouse pointer is positioned over them.

Active alert statuses relevant for specific tiles are indicated by red, orange or yellow symbols on those tiles. These indicators show that one or more alerts are active, but not the number of alerts.

Note

The filter used in the Backend view does not influence the Dashboard display.

Figure 23 Dashboard tiles

The following table describes the tiles displayed on the Dashboard.

Table 27 Dashboard tiles

Tile	Description
1 - Physical Capacity	Displays the raw capacity of the system, rounded off to multiples of 8 GB. The available raw capacity is represented by concentric rings outwards from the center: **Outer Ring** Displays the storage usage using the color codes described in the legend. The actual values are written next to the colored segments. The icon in the bottom right corner of the tile displays the color code legend. The outer ring is divided into colors as follows: • *In use*—From the total, in clockwise direction: shows healthy, degraded and failed capacity. The total of these three items is equal to the sum of capacity in use.

Table 27 Dashboard tiles (continued)

Tile	Description
	• *Unused*—Shows how much further the raw capacity can be expanded; if this capacity is not accessible, it will be marked as Unavailable Unused. • *Special purpose*—Shows spare capacity reserved for system operation; Decreased capacity that was deducted from devices (using the `Set Device Capacity Limit` command), and cannot be used. **Note** A miniature version of the outer Capacity ring is shown in some Backend table views and Property Sheets. **Center circle** Displays the total amount of available raw storage. **Note** Total available raw storage does not represent the total amount of capacity available for volume allocation. **Inner Ring** Displays the snapshot usage. The arc displays the total amount of available data. The filled (bronze) part represents the capacity used by original data volumes, and the hollow (outlined) part represents the capacity used for snapshot volumes. This displays the ratio of snapshot usage. To get a more accurate idea of snapshot usage, see the Backend Capacity Usage view.
2 - I/O Workload	Displays the performance statistics of the system (IOPS, bandwidth and I/O size). More details about I/O can be viewed in the Backend table views: **Application I/O**, **Overall I/O**, and **I/O Bandwidth**. In Advanced Dashboard view (controlled by system Preferences—see Customizing system preferences on page 235) aggregated values of bandwidth and IOPS are displayed. The table in this tile summarizes the Reads, Writes and Totals of IOPS, and throughput and the average size of an I/O.
3 - Internal I/O	The tile displays the system's internal I/O. Clicking on the arrow icon cycles through total internal I/O, internal I/O due to migration, and internal I/O due to rebalance operations.
4 - Rebuild	Indicates if VxFlex OS is currently rebuilding RAID 1 data. A rebuild is usually a result of a recovery due to failure of a server or a storage device. The tile displays the rate in which the data is rebuilt, using a large orange font and icon. Capacity that is still pending rebuild is displayed in small white fonts. In Advanced Dashboard view (controlled by system User Preferences—see Customizing system preferences on page 235) more details are displayed on this tile. Hovering over the title displays a tooltip listing the events that triggered the rebuild operation. Click **More details...** to open the **Monitor** view.
5 - Alert Indicators	Displays the number of active alerts in the system, using the system-wide color codes.
6 - SDCs	Displays the number of SDCs (clients) in the system. The large number in the center is the number of SDCs connected to the MDM. The defined number includes all SDCs defined in the system (some of which may be disconnected from the MDM).

Table 27 Dashboard tiles (continued)

Tile	Description
7 - Volumes	Displays the number of volumes defined across the system, the free available capacity, and the used capacity. The amount of Free capacity shown on this tile is the maximum amount that can be used for creating a new volume. This amount takes into account how much raw data is needed for maintaining RAID 1 and system spares. Note that the number of volumes and the total capacity include snapshots.
8 - Protection Domains	Displays the number and status of all Protection Domains defined in the system. The large number in the center is the number of Protection Domains. This tile is displayed when the dashboard is filtering information according to cluster.
8 - Storage Pools	Displays the number and status of all Storage Pools defined in the Protection Domain. The large number in the center is the number of Storage Pools. This tile is displayed when the dashboard is filtering information according to Protection Domain.
9 - SDSs	Displays the number and status of all SDSs (servers) in the system. The large number in the center is the number of SDSs defined in the MDM. This tile is displayed when the dashboard is filtering according to cluster or Protection Domain. If any SDSs are currently in Maintenance Mode, the orange maintenance icon is displayed on this tile. The specific SDSs currently in Maintenance Mode can be identified using the Backend and Alerts views.
9 - Devices	Displays the number and status of all storage devices defined in the Storage Pool. The large number in the center is the number of devices defined in the MDM. This tile is displayed when the dashboard is filtering according to Storage Pool.
10 - Management	Displays the status of the MDM cluster, or of an MDM operating in Single Mode (one node). The status is displayed graphically as a combination of the MDM cluster elements, and an alert icon if active alerts exist. For more information, see Management (MDM) cluster status on page 155. When you hover your mouse pointer over this tile, a tooltip displays the IP addresses, including the Virtual IP address, used by the MDM cluster or node.
11 - Local Clock	Displays the time on your local VxFlex OS GUI client machine (not the MDM time)

Management (MDM) cluster status

The graphics shown in the following tables represent various MDM states. A tooltip displays the IP addresses used by the MDM cluster or node.

- In 3-node cluster mode, a Master MDM, a Slave MDM, and a Tie Breaker MDM are configured in the system, and statuses are displayed on the **Management** tile of the dashboard.
- In 5-node cluster mode, a Master MDM, two Slave MDMs, and two Tie Breaker MDMs are configured in the system, and statuses are displayed on the **Management** tile of the dashboard.
- In single mode, one Master MDM is configured in the system, and the status of that MDM is displayed on the **Management** tile of the dashboard.

Note

It is not recommended to use Single Mode in production systems, except in temporary situations. The MDM contains all the metadata required for system operation. Single Mode has no protection, and exposes the system to a single point of failure. If the connection to the MDM is lost, the Dashboard is dimmed, and a dialog box is displayed.

The following figure shows how management nodes are displayed on the

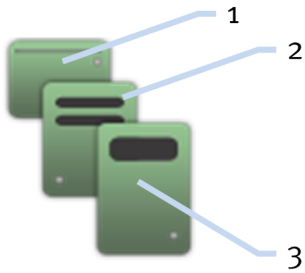

Dashboard.

Legend: 1—Tie Breaker, 2—Slave, 3—Master

The following tables show how management clusters are displayed on the Dashboard.

Table 28 Management node icons with normal operational status (green)

Icon	Description
	5-node cluster
	3-node cluster
	single node

Table 29 Management node status indications and color codes

Example	Description
	Green—Normal operation
	Gold—Degraded state: data is not consistent; system is synchronizing (the node is still operational)
	Gray—Degraded state: a Slave or Tie-Breaker is down
	Blue arrow—An upgrade is in progress
	There is no communication with the Master MDM

Frontend views

The **Frontend** view provides detailed information about frontend objects in the system, including volumes, SDCs and snapshots, and lets you perform various configuration operations. The main areas of the Frontend view are:

- **Filter**—lets you filter the information displayed in the table and Property Sheets. For more information, see Filter (1) on page 163.

- **Toolbar buttons**—let you perform commands on the selected row in the table (add, remove, configuration), or toggle the display of Property Sheets, by clicking the appropriate button
- **Table**—displays detailed information about system objects. The table displays a wide range of information, which can be filtered. Certain commands can be performed on objects, using the context-sensitive menu for the desired row in the table, or the **Command menu** on the toolbar.
- **Property Sheets**—display very detailed read-only information about the object selected in the table. For more information, see Property Sheets on page 168. For more information about the terminology used in the Property Sheets, see the Glossary at the end of this publication.

Volumes

Figure 24 Frontend > Volumes view

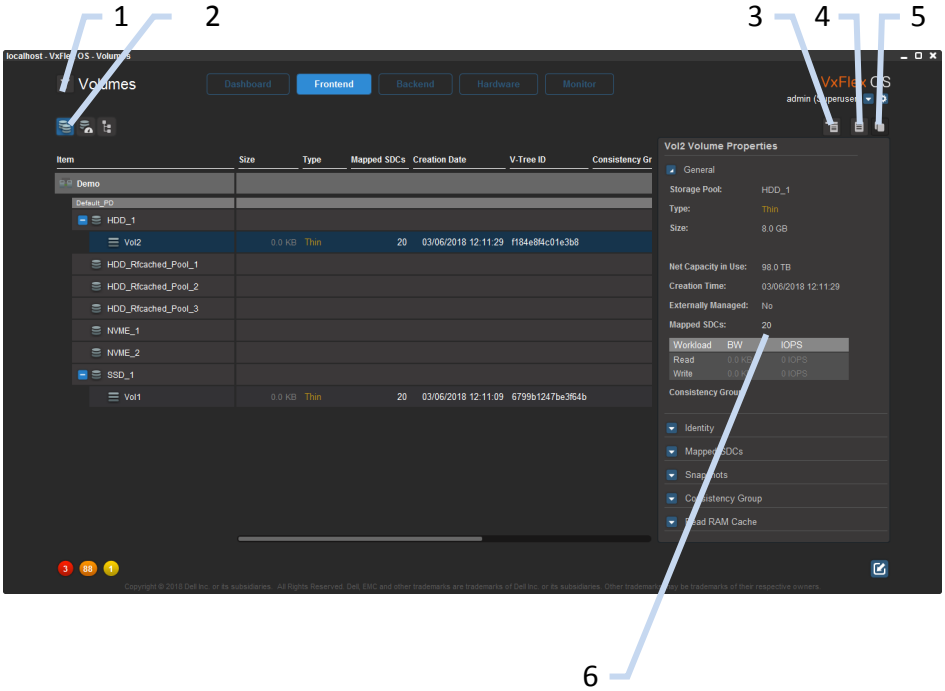

1 **Filter toggle**—controls display of the filter

2 **Table view options**—each button provides a different combination of properties which can be displayed together in the table.

3 **Command menu**—contains a list of commands which you can perform on the row selected in the table

4 **Show Property Sheet**—controls display of the Property Sheet. The Property Sheet displays information about the object selected in the table.

5 **Duplicate Property Sheet**—Opens and floats the Property Sheet in a new window, which can be kept open while you display the **Table**properties of a different object, for comparison or other purposes

6 **Property Sheet**—The Property Sheets provide detailed read-only information about the object selected in the table

SDCs

Figure 25 Frontend > SDCs view

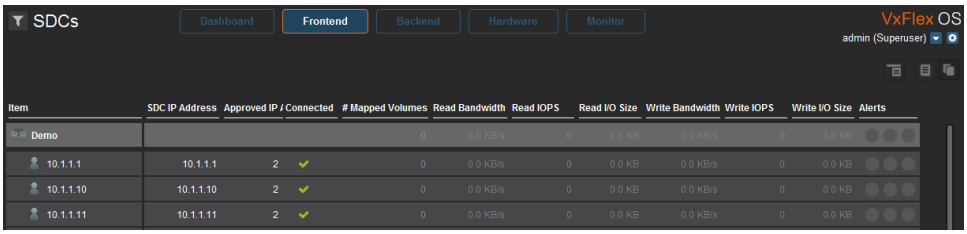

Snapshots

Figure 26 Frontend > filtered Snapshots view

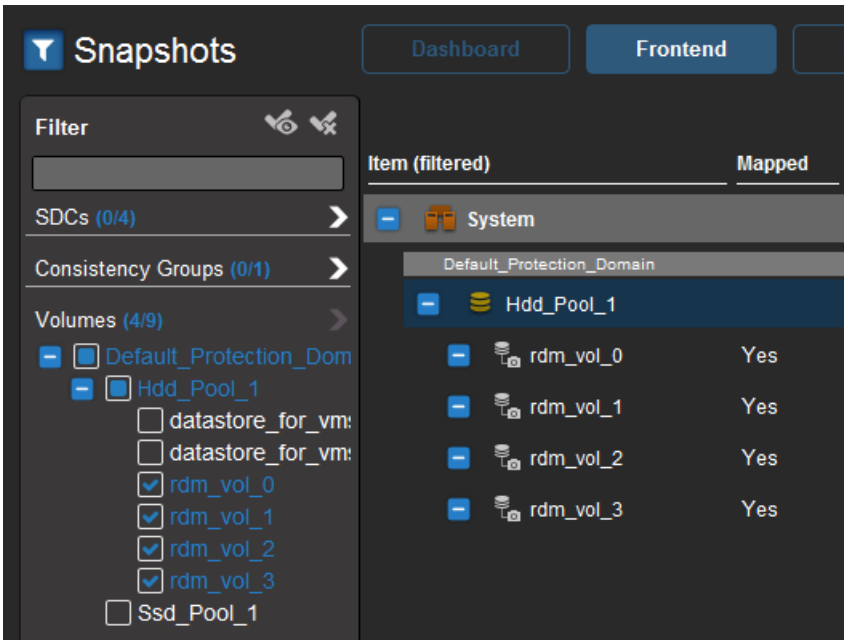

Adding Frontend objects to the filter

This topic explains how to add Frontend objects such as volumes, SDCs, snapshot trees and consistency groups to the Frontend filter. The filter simplifies the display in the main window, by showing only required objects, and hiding the rest.

Note

Toggle the filter display on and off by clicking the filter icon at the top left corner of the window. When the filter is on, only items added to the filter will be visible in the table.

To add objects to the filter, perform these steps:

Procedure

1. In the **Frontend** view, navigate to the objects that you want to add to the filter, and select them.
2. From the **Command menu** or context-sensitive menu, select one of the following, according to your needs:
   - **Add to Filter**
   - **Add Snapshot Tree to Filter** (adds all the snapshots for the selected volume)
   - **Add Consistency Group to Filter** (adds all the members of the consistency group)

Compute

Figure 27 Frontend > Compute view

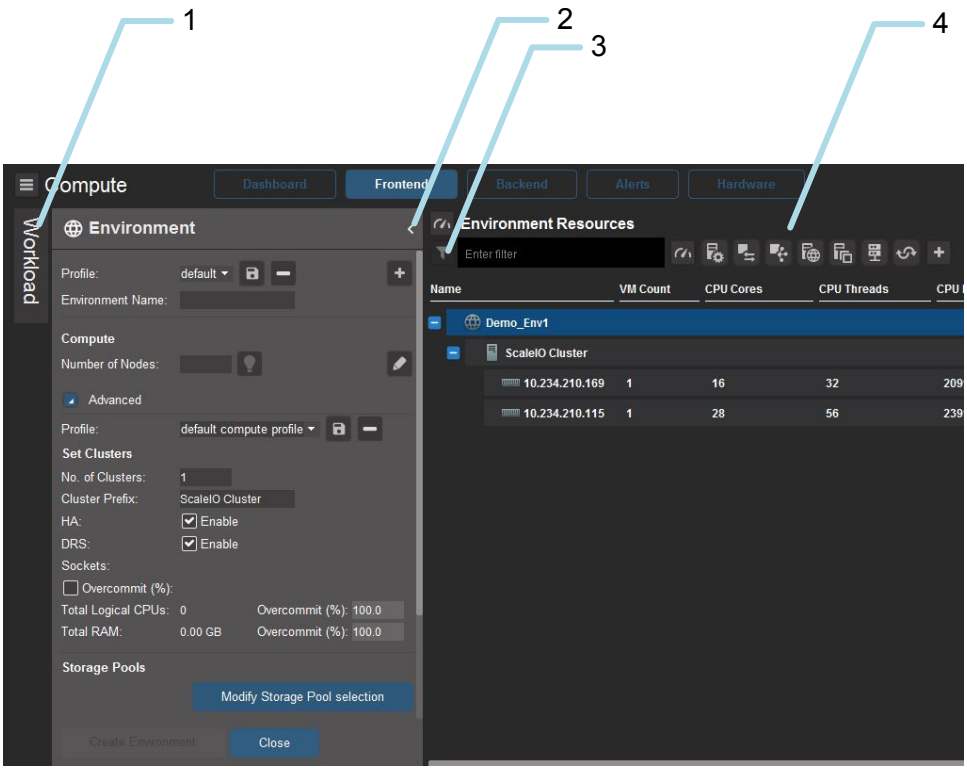

1 **Environment** and **Workload** sidebars—display the **Environment** and **Workload** panes

2 **Hide arrow**—hides the **Environment** or **Workload** pane

3 **Filter** button and box—controls display of the information shown in the table. To filter using free text, type text in the box.

4 **Toolbar**—buttons control the information displayed in the table. The **More Details** button displays the available commands in a drop-down menu. The **+** button lets you create a new environment. The remaining buttons can be used for

monitoring purposes, and are described in Monitoring compute and network environments on page 176.

Backend view

The **Backend** view provides detailed information about backend objects in the system, and lets you perform various configuration operations. The main areas of the Backend view are:

- **Filter**—lets you filter the information displayed in the table and Property Sheets. For more information, see Filter (1) on page 163.
- **Toolbar buttons**—let you display various sets of information, perform commands on the selected row in the table (add, remove, configuration), or toggle the display of Property Sheets, by clicking the appropriate button
- **Table**—displays detailed information about system objects. The table displays a wide range of information, which can be filtered. Certain commands can be performed on objects, using the context-sensitive menu for the desired row in the table, or the **Command menu** on the toolbar.
- **Property Sheets**—display very detailed read-only information about the object selected in the table. For more information, see Property Sheets on page 168. For more information about the terminology used in the Property Sheets, see the Glossary at the end of this publication.

Note

Irrelevant, mostly zero values, are “dimmed” in the Backend view. The only exception (where a value could be greater than zero but is dimmed) is when the Max Capacity is the same as the Total Capacity. In this case, Max Capacity would be dimmed (if decreased capacity exists, Max = [Total + decreased] in the Backend table). Similarly, arrows and icons are unavailable in the same circumstances.

Note

Some objects in the system can be identified by ID numbers. The ID numbers displayed in the VxFlex OS GUI can be used in CLI commands to specify these objects.

Figure 28 Backend view

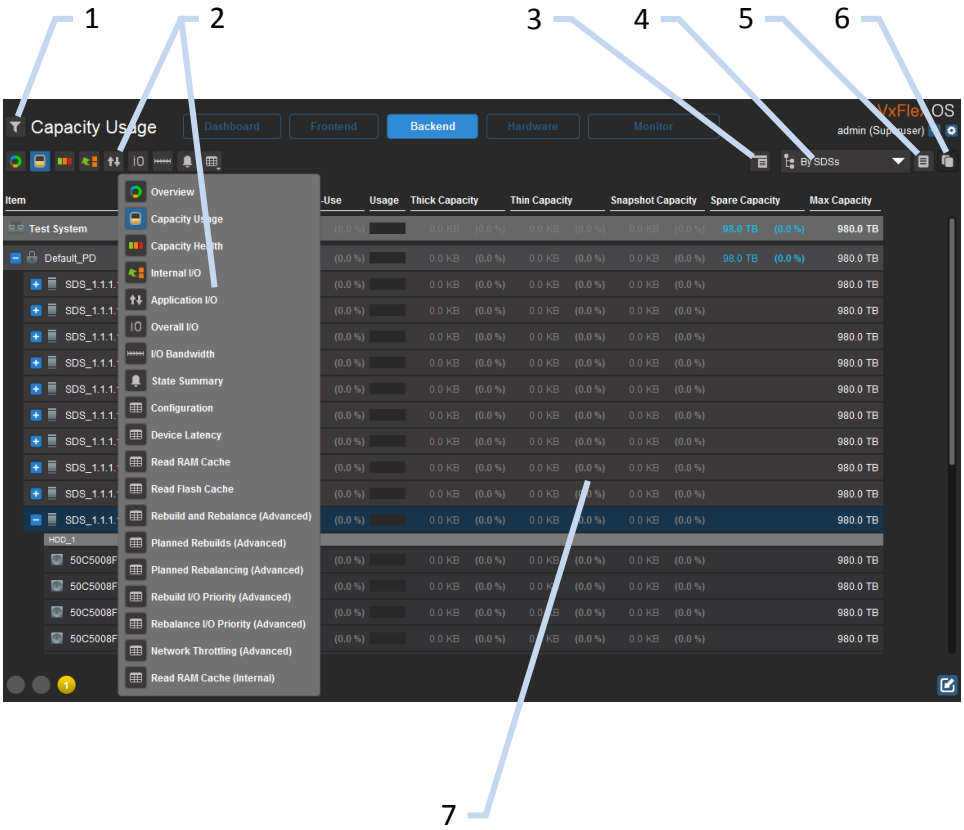

1 Filter toggle—controls display of the filter

2 Table view options—each button provides a different combination of properties which can be displayed together in the table. Additional views are available from the **More Table Views** button. For more information, see Table on page 163.

3 **Command menu**—contains a list of commands which you can perform on the row selected in the table. For more information, see Command menu (3) on page 163.

4 **Storage Pools\SDS** toggle—toggles display of the table rows grouped according to either Storage Pools, or SDSs. Some table views are only available when sorted by either SDSs or Storage Pools. For example, Fault Sets are only displayed when

5 **Show Property Sheet**—controls display of the Property Sheet. The Property Sheet displays information about the object selected in the table.

6 **Duplicate Property Sheet**—Opens and floats the Property Sheet in a new window, which can be kept open while you display the properties of a different object, for comparison or other purposes

7 **Table**—displays a summary of properties for the objects selected in the filter, according to the selected table view option

sorting by SDSs, and Rebuild I/O Priority is only displayed when sorting by Storage Pools.

Command menu (3)

The **Command menu** button displays a list of commands which you can perform on rows selected in the table. The contents of the **Command menu** differ, depending on the object selected in the table. Many of the commands can also be accessed from the context-sensitive menu when table rows are right-clicked.

Filter (1)

The filter lets you filter out the information shown in the table, so that only the objects related to the items selected in the filter are visible. Items colored dark blue in the filter are displayed. The buttons at the top of the filter panel control the items hidden, displayed, selected or cleared in the filter. Below these buttons, a search field lets you type free text to search for a object by name. The filter below shows a filter being used to show specific devices in the Backend table. In the Backend, you can filter information according to Storage Pools, SDSs, Volumes, Protection Domains and Devices.

Your filtering choice would depend on your current problem or management needs. You can also add a specific object to the filter, using the `Add to Filter` command from the **Command menu** or context-sensitive menu.

Figure 29 Backend filter

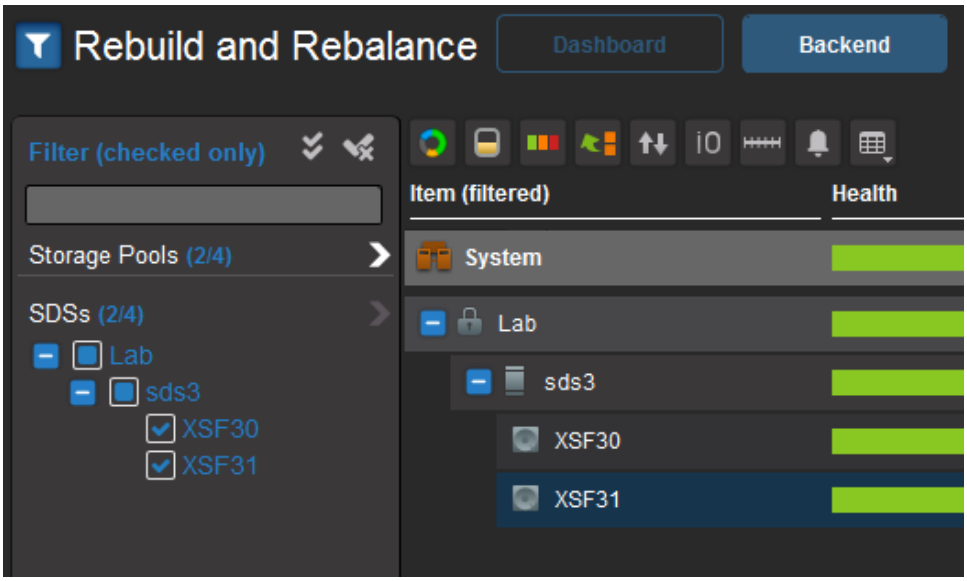

Table

The table displays object statuses, and allows configuration of some objects. The contents of a Property Sheet for a specific object are determined by the row selected in the table. The color codes described in Color codes on page 147 also apply to the items displayed in the table.

The information displayed in the table is controlled by the filter and the table view options. Selected objects (rows in the table) can be created, configured and removed

using context-sensitive menu options, or commands from the **Command** menu. The columns in the table can be resized, by dragging the borders in the heading row of the table. The scrollbar at the bottom of the table lets you scroll through the columns in the table. The following table describes the available table views:

Note

Clock icons in the cells indicate stale\aging data.

Table 30 Table view options

Name of view	Contents of view	Suggested use, comments
Overview	Total Capacity, Capacity In-Use, I/O Bandwidth, IOPS, Rebuild, Rebalance, Alerts	Provides a general overview the capacity and health of system objects. If you find that there are active alerts, you can switch to **State Summary** table view, or the **Alerts** view to see more details about the alerts. While this view is similar to the Dashboard, the main difference is that here, you can simultaneously monitor many objects within their hierarchy. You can then filter the table for specific objects or sets of objects, in contrast to the Dashboard, where you can only drill-down to objects. In addition, commands and Property Sheets are available to you from this view. **Note** The miniature Capacity ring shown here represents the outer Capacity ring on the Dashboard. For more information, see Dashboard tiles on page 152.
Capacity Usage	Total, In-Use, Usage, Thick Thin, Snapshot, Spare, Max Related Property Sheet section: **Capacity**	Provides a breakdown of capacity usage per use type. You can use this to check whether more capacity needs to be added to your system, and where.
Capacity Health	Capacity In-Use, Protected, Degraded, Failed, Health, Rebuild, Rebalance Related Property Sheet sections: **Capacity**, **Alerts**, **Rebuild/Rebalance**	Provides information about the health of the objects in the system, per object
Internal I/O	Health, Backward Rebuild, Forward Rebuild, Rebalance Related Property Sheet sections: **Alerts**, **Rebuild/Rebalance**	Provides a summary of Rebuild and Rebalance health, status, and workload per object
Application I/O	Bandwidth, IOPS and I/O Size for: Total, Read, Write, 2nd Write Related Property Sheet section: **Workload**	Provides workload information for applications reading/writing to storage in the system. 2nd Writes refer to the protection copy of data being written to storage.

Table 30 Table view options (continued)

Name of view	Contents of view	Suggested use, comments
Overall I/O	Bandwidth, IOPS and I/O Size for: Total, Total Read, Total Write Related Property Sheet section: **Workload**	Provides workload information for all I/Os in the system, including both application I/Os, and I/Os for internal processes
I/O Bandwidth	Read, Write, Backward Rebuild, Forward Rebuild, Rebalance, Total, Total Read, Total Write Related Property Sheet section: **Network Throttling**	Shows the bandwidth being used for various jobs in the system
State Summary	Summary Related Property Sheet section: **Alerts**	Can be used to identify items in the system which have open alert states, and to view the alert messages, including information statuses of various objects. If there are pending security certificates, and if an SDS is in Maintenance Mode, this is also indicated here. Alerts marked blue are for information purposes only, and do not require that any action be taken.
Configuration	Total Capacity, SDSs, Devices, Storage Pools, Volumes, Free Capacity for Volume Allocation, Alerts Related Property Sheet sections: **Identity**, **Related Objects**	Provides an overview of the number of objects per type in your system, their capacity, and lets you determine the amount of free capacity available for creating an additional volume. The **Free Capacity for Volume Allocation** column is the only one that is not stated in Raw Capacity values. This table view is a convenient location from which to perform Add, Remove, Activate and Inactivate commands. For more information, see Configuring the System using the VxFlex OS GUI on page 179. The **Related Objects** section of the Property Sheet helps you to identify related objects to the one selected in the table.
Device Latency	Average Read Latency, Average Write latency, Average Read Size, Average Write Size, Scanned Capacity, Resolved Errors, Data Conflicts Related Property Sheet sections: **Device Latency** , **Background Device Scanner**	Provides an overview of performance information, and Background Device Scanner results, per device. When the table is sorted by SDSs, the Down arrow in each SDS row reveals all the devices in the SDS. **Note** Information is shown only at Device level, and there is no aggregation of information at higher levels. If the background device scanner is enabled, several device read statistics are dramatically affected.
Read RAM Cache	State, Size, Used, Hit Rate, Write Mode Related Property Sheet section: **Read RAM Cache**	Lets you check which SDSs have Read RAM Cache enabled, view associated counters, and check which Storage Pool is set to use the cache

Table 30 Table view options (continued)

Name of view	Contents of view	Suggested use, comments
		for its devices. Advanced feature; modify with caution.
Read Flash Cache	State, Size, Used, Read Hit Rate, Total Errors, Alerts Related Property Sheet sections: **Read Flash Cache** and **Related Objects**	Lets you check which SDSs and Storage Pools have Read Flash Cache enabled, view associated counters. The **Related Objects** section of the corresponding Property Sheet shows the names/paths of the devices used for caching.
Rebuild and Rebalance (Advanced)	Health, Backward Rebuild, Forward Rebuild, Rebalance Related Property Sheet section: **Rebuild/Rebalance**	Shows workload and active rebuild and rebalance jobs, as well as pending jobs to be processed. Credited incoming and outgoing information is displayed for each process (Backward Rebuild, Forward Rebuild and Rebalance).
Planned Rebuilds (Advanced)	Degraded, Distribution, Backward and Forward Rebuilds Related Property Sheet section: **Rebuild/Rebalance**	Shows amount of unprotected data, a visual breakdown of rebuild progress (Distribution), bandwidth, direction and status (Active/Pending) of jobs. You can monitor how the degraded capacity is planned to be rebuilt: How failed degraded capacity is expected to be rebuilt, then how the system will pool it for the rebuild processes (backward and forward,) and which of it is actively being rebuilt. **Note** Overall Degraded capacity includes both healthy and failed copies of the data, while the processed capacity includes only the failed copies. Advanced feature; modify with caution.
Planned Rebalancing (Advanced)	Protected, Distribution, Rebalance Related Property Sheet section: **Rebuild/Rebalance**	Shows amount of protected data, a visual breakdown of rebalance progress (Distribution), bandwidth, direction and status (Active/Pending) of jobs. Advanced feature; modify with caution.
Rebuild I/O Priority (Advanced)	Policy, Concurrent I/O limit, Bandwidth Limit, Application Threshold, Quiet Period Related Property Sheet section: **I/O Priority**	Displays bandwidth settings currently configured for Rebuild jobs. These I/O Priority settings apply only to Storage Pools. These settings control system performance. Advanced feature; modify with caution.
Rebalance I/O Priority (Advanced)	Policy, Concurrent I/O limit, Bandwidth Limit, Application Threshold, Quiet Period Related Property Sheet section: **I/O Priority**	Displays I/O priority settings currently configured for Rebalance jobs. These I/O Priority settings apply only to Storage Pools. These settings control system performance. Advanced feature; modify with caution.
Network Throttling (Advanced)	Overall I/O Limit, Rebuild I/O Limit, Rebalance I/O Limit, Rebuild Incoming Limit, Rebalance Incoming Limit, Rebuild Queue Length,	Displays Network Throttling settings currently configured in the system. These settings control

Table 30 Table view options (continued)

Name of view	Contents of view	Suggested use, comments
	Rebalance Queue Length, Rebuild Outgoing Jobs, Rebalance Outgoing Jobs Related Property Sheet section: **Network Throttling**	system performance. Advanced feature; modify with caution.
Read RAM Cache (Internal)		For support purposes only. Visible only if Backend Internals are enabled in the **User Preferences** window.

Alerts view

The **Alerts** view provides a list of the alert messages currently active in the system, in table format. You can filter the table rows according to alert severity, and according to object types in the system. For a list of alerts generated by the system, see VxFlex OS Alerts in SNMP, VxFlex OS GUI, REST, and ESRS on page 283

To view a Property Sheet for a specific alert, select the corresponding row in the table, and click the **Show Property Sheet** button. For more information about Property Sheets, see Property Sheets on page 168.

Figure 30 Alerts view

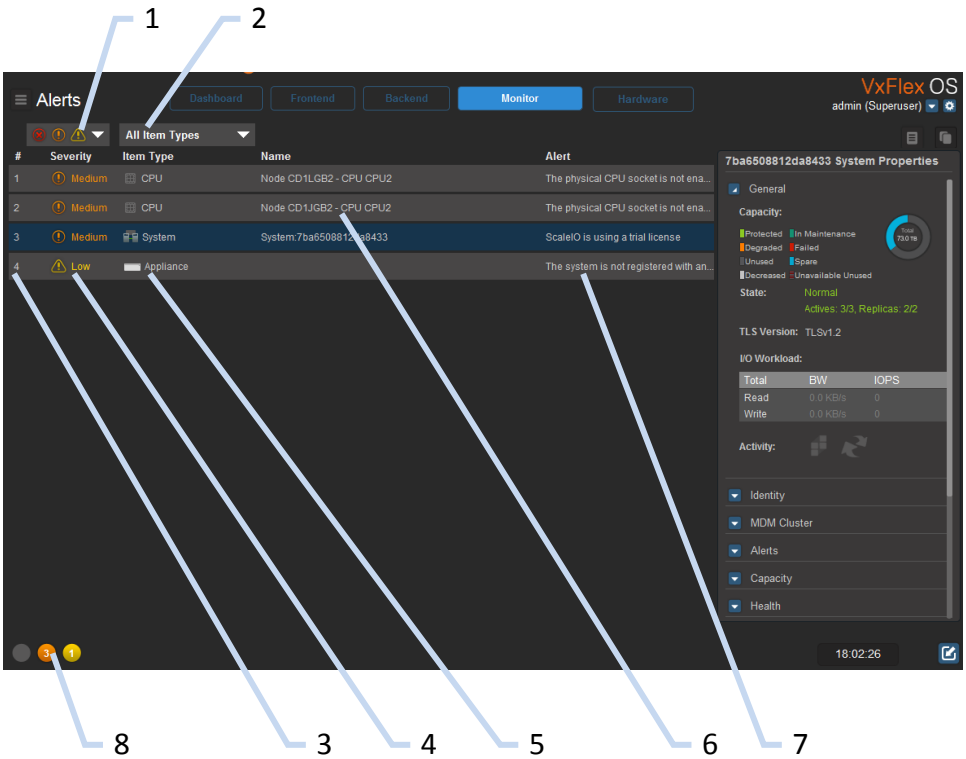

1 Severity filter—filters the table contents according to alert severity:

5 **Item Type**—MDM, Protection Domain, Storage Pool, SDS, Device, SDC

	All alerts **Medium**—alerts which are Medium or High severity **High**—only alerts which are High severity		
2	Item Types filter—filters the table contents according to: All Item Types, MDM, Protection Domain, Storage Pool, SDS, Device, SDC	6	**Name**—the user-defined name of the item, if one has been defined
3	**#**—the line number in the Alerts table	7	**Alert**—the alert message generated by the system
4	**Severity**—the alert severity: High, Medium or Low	8	Alert indicators—summarize the total amount of each alert type (from left to right: High, Medium, Low)

Property Sheets

The Property Sheets provide detailed read-only information about the object selected in most **Frontend** tables, **Backend** tables, **Monitor** tables.

In addition, the **Hardware** view displays detailed information about selected hardware components.

The contents of the Property Sheets differ, depending on the object selected in the table. Property sheets help you to monitor specific objects in the system by displaying the following, using the blue collapse and expand arrows next to each section of the Property Sheet:

- **General** and **Health** information about the object
- **Identity**—identifying information, such as an object's ID number, name, IP addresses, port usage, VM usage, GUID
- **Mapped Volumes** for the selected SDC, including the volume name, and bandwidth limit and IOPS limit per volume
- **MDM Cluster** details, such as cluster mode and state, IP addresses and ports (including the Virtual IP address) used for management and MDM functions, and SSL version being used for secure communication. The Virtual IP addresses and the Virtual IP interfaces are displayed together in table format, to indicate the mapping between them.
- **Alerts**—alert status per selected object, including SDS Maintenance Mode (if active)
- **Capacity** usage per selected object
- **Workload** information (bandwidth usage and IOPS) per selected object
- **Rebuild/Rebalance** information about the selected object(s), for forward and backward rebuilds, rebalancing, data at rest, and job status (active\pending). The type of information displayed depends on the type of object selected in the table.
- **Device Latency** averages for Read and Write, and average I/O size, for the device selected in the table

Note

If the background device scanner is enabled, several device read statistics are dramatically affected.

- **Read RAM Cache** configuration, state, and statistics for the selected object. The type of information displayed depends on the type of object selected in the table.
- **Read Flash Cache** configuration, state, and statistics for the selected object. The type of information displayed depends on the type of object selected in the table.
- **Device Test Results** for the device selected in the table, if any tests have been performed
- **Background Device Scanner** results for the device selected in the table, if the scanner is enabled
- **Network Throttling** configuration for the selected SDS, including bandwidth limit per job, and queue length, for both Rebuilding and Rebalancing.
- **I/O Priority** configured for the selected Storage Pool, including Rebuild and Rebalance states, number of parallel jobs, I/O prioritization policy, concurrent I/Os and bandwidth limit
- Miscellaneous items, such as DRL mode, zero padding, checksum mode
- **Performance** profile currently assigned to the selected object: Default, High or Custom. “Custom” is displayed whenever performance-related parameters have been configured manually, instead of via a performance profile.
- **Oscillating Failure Counters** are shown for SDSs and devices selected in the Backend view, for SDCs selected in the Frontend view, and for Alert messages.
- **Oscillating Failure Parameters** currently configured in the system are shown for the entire system, for the selected Storage Pool or Protection Domain, and for Alert messages. The counters shown depend on the object selected in the table.
  - **Window:** the sliding time-window for each interval (Short, Medium and Long)
  - **Threshold:** the number of errors that may occur before error reporting commences
  - **Period:** the time interval of each Window, in seconds
- **Maintenance Mode** state of the selected SDS is shown here. States include: No Maintenance and In Maintenance Mode.
- **Certificate Info** is shown for security certificates
- **Related objects**, which can be very useful for troubleshooting problems, or for planning purposes, when you need to make changes to your system.
- From the **Hardware** view, you can display property sheets for various hardware components (when present, depending on server model) by clicking on the corresponding buttons or graphical representations in the middle panel of the window. By default, when you toggle between nodes, the Property Sheet initially displays Node Properties.
  - **BMC (iDRAC)** (Baseboard Management Controller) displays connectivity status, addresses and firmware version of the node’s BMC (iDRAC)
  - **Fan** displays the status and speed (RPM) of each fan in the node
  - **BOSS card** displays general inventory information and physical information.
  - **NIC** (Network Interface Card) lets you toggle between NICs in the node. Information is displayed for each built-in NIC’s PCI slot and ports.

- **CPU** (Central Processing Unit) displays details about the node's CPUs, including inventory and specifications, temperature, voltage and cache information for L1, L2 and L3
- **RAM** displays the node's RAM capacity, inventory, size, speed, temperature and number of errors according to each DIMM slot. NVDIMM Battery Properties shows the health and state of the NVDIMM's onboard battery.
- **Controller** shows RAID storage controller properties, such as inventory, temperature, specifications, IDs, battery, and bus location in the node
- **PSU** (power supply unit) displays the state and input of the node's PSUs
- **Node** displays inventory details, IP addresses and roles, DNS and NTP IP addresses, boot disk information, temperature and voltage readings at a number of physical locations in the node, and air flow.
- **Cache** displays information about the caching used in the node, such as cache size, a list of the devices being used to cache data, and a list of the HDDs being cached
- **Physical drives (devices)** each display location, inventory, specifications, temperature and other health statuses, endurance (where relevant), configuration, and logical drive properties. Acceleration properties for the selected device are also shown here, in the Logical Drive Properties section of this Property Sheet.

You can view properties for multiple objects by using the **Duplicate Property Sheet** button, and then navigating to a different object's row in the table. When more than one Property Sheet is open, a floating widget that controls them is displayed in the bottom part of the main window, as shown in the following figure.

Note

You can open Property Sheets, duplicate them, and then simultaneously work on other unrelated objects in the system. The duplication feature is not supported in the Hardware view.

Figure 31 Multiple floating Property Sheets

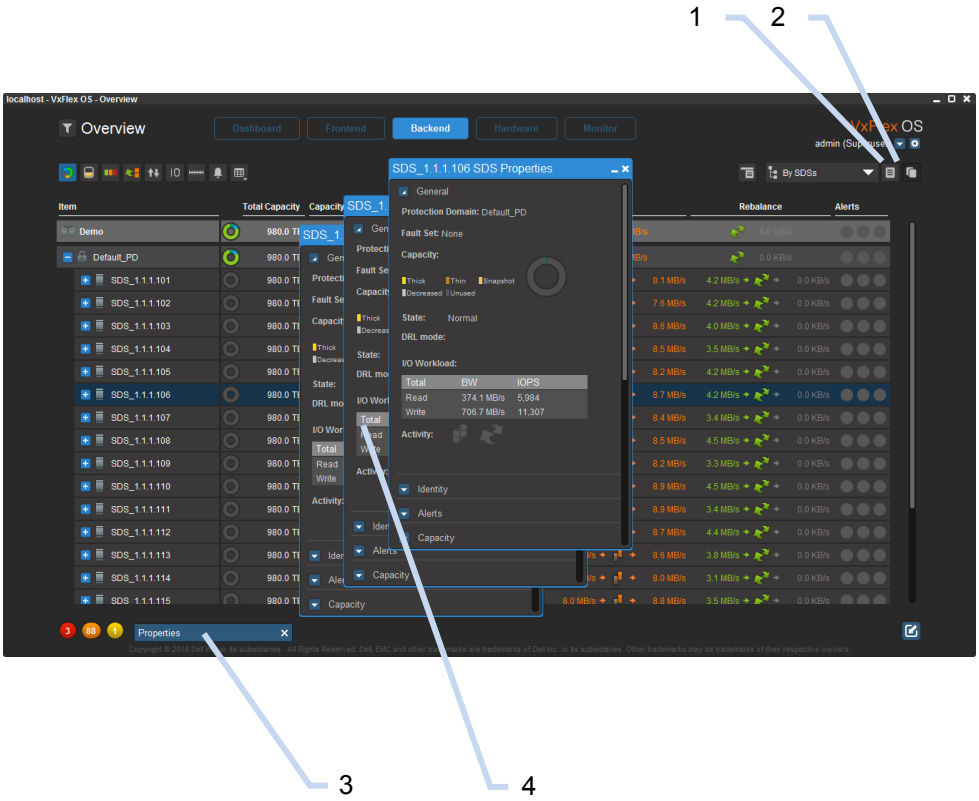

1 **Display\Hide Property Sheet** button

2 **Duplicate Property Sheet** button

3 **Duplicated Property Sheets** toggle

4 Multiple Property Sheets opened using the **Duplicate Property Sheet** button

The following figure shows a typical Property Sheet for an SDS, with the **Identity** section of the Property Sheet expanded.

Figure 32 Example of a Property Sheet for an SDS

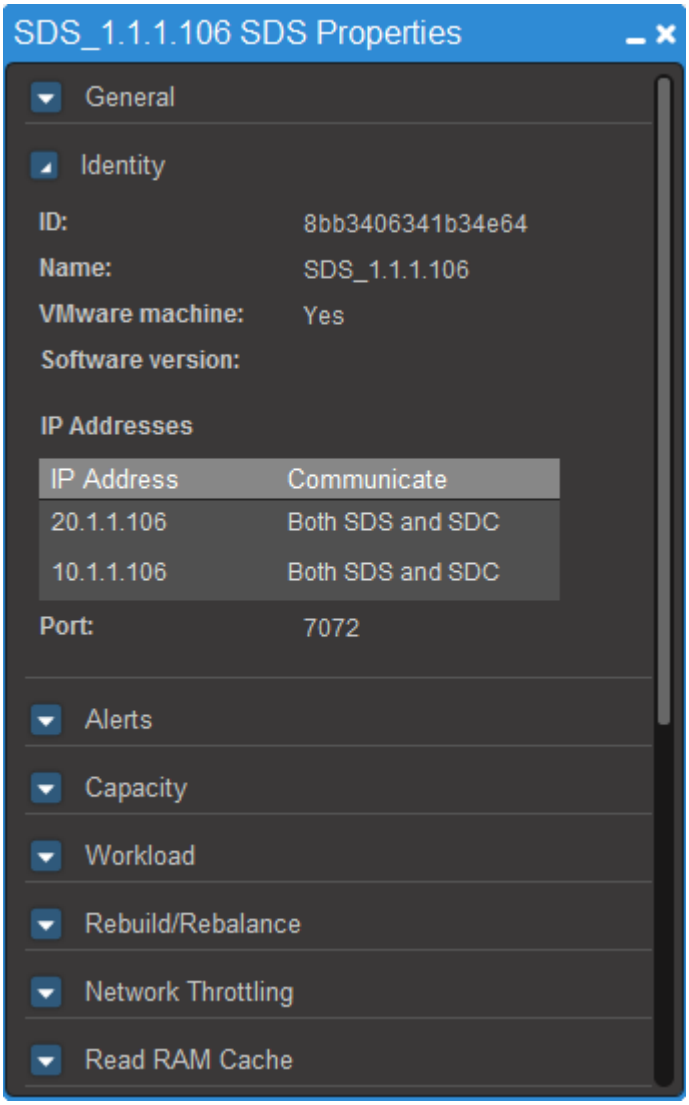

CHAPTER 9

Monitoring the System using the VxFlex OS GUI

The following topics explain how to monitor the system, using the VxFlex OS GUI.

- Viewing object properties in Backend and Frontend views 174
- Viewing detailed hardware properties 174
- Monitoring compute and network environments 176
- Viewing licensing information 177
- Verifying your connection to the Management cluster 178
- Collecting logs 178

Viewing object properties in Backend and Frontend views

Procedure

1. In **Backend** or **Frontend** view, use the filter to display one or more objects, and select the corresponding check boxes of the desired objects (optional).
2. Navigate to the desired object in the table.
3. Display the required information, using the table view options.
4. Select the required object's row in the table, and then, on the expandable Property Sheet on the right side of the window, click the blue arrow buttons beside the headings to expand them and view specific status information.

Note

Contents of the Property Sheet are dynamic, and differ, depending on the row selected in the table. For more information about Property Sheets, see Property Sheets on page 168.

Viewing detailed hardware properties

This section explains how to view detailed information about your system's hardware components.

Identifying hardware errors

There are two ways of identifying hardware errors in the system:

- From the **Alerts** view—alerts for hardware components are displayed in the Alerts table
- From the **Hardware** view—components with error states are shown with a colored outline. The error state colors are uniform throughout the VxFlex OS system. An error state outline is displayed around the corresponding button or component in the middle panel of the **Hardware** view.

Note

Blue is used to indicate that a component is selected. It does not indicate an error state.

Identifying the IP role of a node

Every VxFlex OS node is configured with both SDS and SDC roles.

Identifying the SDS running on a node

The IP address of the SDS running on a node can be identified using the Hardware view Property Sheet.

To view identify an SDS, perform the following:

Procedure

1. In the **Hardware** view, navigate to the node, and in the middle panel, click the **Node** button.

The Property Sheet displays information about the node. The **SVM Management IP** address shown on the Property Sheet is the SDS's IP address.

Identifying the SSD used for caching

The Property Sheet in the Hardware view can be used to determine whether the SSD is being used for caching.

To view SSD use, perform the following steps:

Procedure

1. In the **Hardware** view, navigate to the node, and in the middle panel, **Front** view click one of the storage devices.

   The information displayed differs depending on the selected node and device.

   Figure 33 A selected device (1) and its corresponding Property Sheet (2)

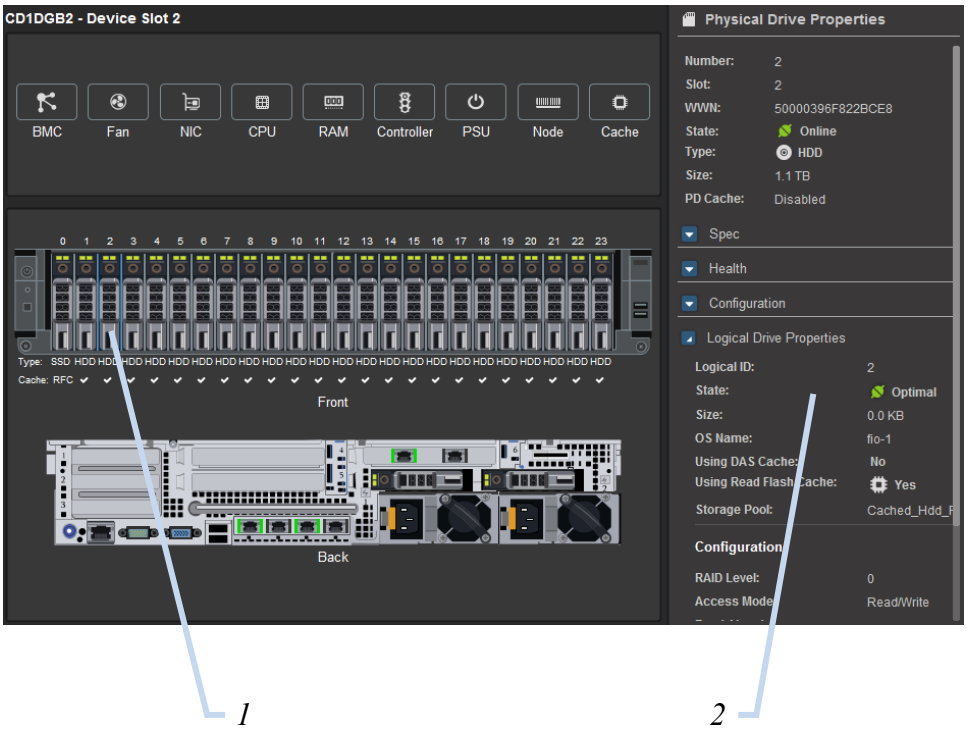

2. In the uppermost part of the **Physical Drive Properties** Property Sheet, check if the **Type** is SSD or HDD.

   a. If the drive type is SDD, proceed to the next step.

   b. If the drive type is HDD, click other drives in the **Front** view until you find an SSD, and then proceed to the next step.

3. Expand the **Logical Drive Properties** section. If the SSD drive is used for caching, the property Used as Read Flash Cache: Yes is displayed. If the SSD drive is not used for caching, the property Used as Read Flash Cache: No is displayed.

   a. If the SSD drive is used for caching, the property **Used as Read Flash Cache: Yes** is displayed.

   b. If the SSD drive is not used for caching, the property **Used as Read Flash Cache: No** is displayed.

Figure 34 Examples of drives: type (1, 2), storage (3), and caching (4)

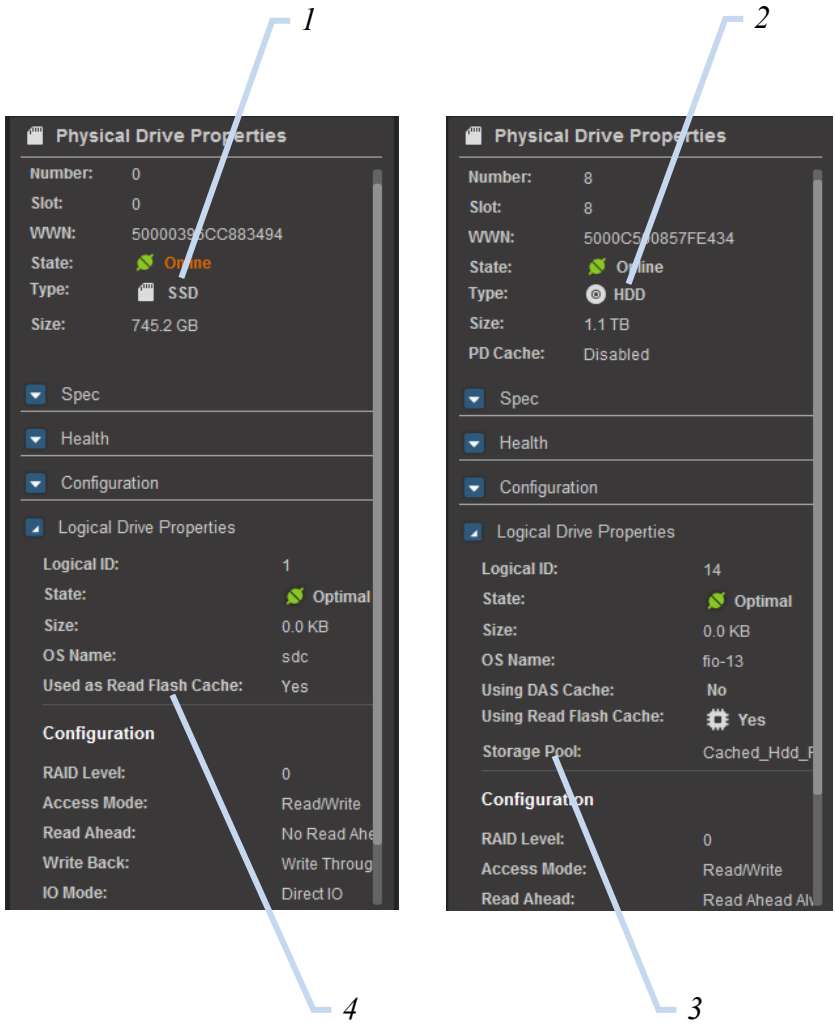

MTU values

The default MTU (Maximum Transmission Unit) values in the system are 1500. If you want to modify these values, refer to the Technical Note *Fine-Tuning VxFlex OS Performance* for more information.

Monitoring compute and network environments

The VxFlex OS GUI provides detailed information and statuses for the compute and network environments, allowing you to monitor them from the VxFlex OS GUI, instead of having to consult the vCenter.

The toolbar buttons in the **Compute** view provide several ways of displaying compute and network statuses and configuration details. The following table summarizes the items that can be monitored.

Compute toolbar button	Displays...
Environment Resources	For each host: the number of VMs, CPU cores, and CPU threads; CPU Mhz and RAM MB

Compute toolbar button	Displays...
Cluster Configuration	For each cluster: HA state (High Availability), DRS state (Distributed Resource Scheduler), Overcommit Enabled, CPU Overcommit Ratio, RAM Overcommit Ratio
vMotion Networks	For each virtual NIC: MTU and VLAN configuration
VM Networks	For each virtual NIC: MTU and VLAN configurations
VMs by Clusters	For each VM (displayed according to cluster > host): VM status, Data Center, vCenter Server, Networks, VM Tools, OS (operating system), Cores, Memory, NICs, Storage size, Datastore names, VM template
VMs by Workloads	For each VM (displayed according to cluster > host): number of VMs per workload, Storage Pools, VM Template, VM Status, Host, Data Center, vCenter Server, Networks, VM Tools, OS (operating system), Cores, Memory, NICs, Storage, Datastores
Datastores	For each datastore (displayed according to Storage Pool): file system type, size, and volume name
Command Monitor	For each process: process name, reported entity, start and end time, status of the process, and error message

Procedure

1. Click the **Frontend** tab, and select **Compute**.

   The **Compute** view appears.
2. Click the required button on the toolbar.
3. Click a row in the table in the main part of the window, and then expand the row to reveal more details.

Viewing licensing information

Before a permanent license is installed, a banner is displayed at the top of the VxFlex OS GUI window that provides links for purchasing a license, and for registering for VxFlex OS newsletters.

Information required for licensing purposes is located in the **About** window, as described in the following steps:

Procedure

1. From the **System Settings** menu at the top right side of the window, in any view, select the **About** option.

The **About** window is displayed.

2. Make a note of the information displayed for **Installation ID**, which is required for electronic licensing purposes.

   Additional information pertaining to your license is also displayed in this window.

Verifying your connection to the Management cluster

If your VxFlex OS GUI connection to the system is operating normally, regular access to all the VxFlex OS GUI views is possible.

You can check which MDM you are connected to, by hovering your mouse over the Management tile on the Dashboard.

The Management IP address to which you are connected is also displayed in the top left corner of the VxFlex OS GUI. For more information about MDM statuses, see Management (MDM) cluster status on page 155.

If the connection is lost, the VxFlex OS GUI is dimmed, and a dialog box is displayed.

Collecting logs

Use the VxFlex OS GUI to collect a downloadable log file in *.ZIP file format from the system.

This operation collects logs from MDM, SDS, AMS, iDRAC, RAID controller, and OS (including ESXi VM-support) logs. Logs are collected in ZIP format, which you can download for viewing.

This procedure is used for technical support purposes only, and is described in the *VxFlex OS Log Collection Guide*.

CHAPTER 10

Configuring the System using the VxFlex OS GUI

The following topics explain how to configure the system, using the VxFlex OS GUI.

- Configuring capacity 180
- Modifying cache settings 191
- Configuring volumes, volume trees, SDCs, and snapshots 196
- Entering and exiting SDS Instant Maintenance Mode 205
- Cluster operations 206
- Configuring Oscillating Failure counters 212
- Applying Performance Profiles to system components 215
- Configuring I/O priorities and bandwidth use (advanced) 216
- Enable and disable Rebuild/Rebalance (advanced) 217
- Using the background device scanner 218
- Modifying Checksum protection mode 219
- Renaming objects 220
- Provision compute, storage and network environments, and deploy workloads 220
- Approving pending security certificates 235
- Customizing system preferences 235
- Modifying system settings 238
- Changing login credentials for existing Management users 243

Configuring capacity

Add, remove, and configure capacity.

The following topics explain how to add, remove, activate, and inactivate capacity, activate devices, clear device errors, and set device capacity limits.

You can also add and remove nodes, add and remove devices, and turn on identification LEDs.

Adding nodes

If you want to extend a system, or replace a node, use the **Add Nodes** command, which is accessed from the **System Settings** menu at the top right corner of the main window. Follow the procedures described in the *VxFlex Ready Node Deployment Guide*, "Deploying the System".

Note

Nodes must be in factory state before adding them to a system. To reset a node to default factory settings, see Perform a factory reset on page 314.

Remove the node from the system

Remove a node from the VxFlex OS system using the AMS.

Before you begin

Note

During this procedure, a rebalance may take place. You may want to plan this procedure for a scheduled maintenance window.

- On ESXi nodes if the faulty server is up, migrate all VMs except for the Storage VM to another server. If the server is not up, use the vCenter client to start these VMs on another server.

  Note

  If VMware HA is enabled, VMs can migrate without service interruption.

- If the node includes an SDC, volumes may be mapped to the node. In this case, during the removal process you will be alerted to first remove the mapped volume. Refer to *VxFlex Ready Node AMS User Guide* for the relevant steps.
- All appliance nodes act as an SDS and possibly an SDC. If the server is also an MDM or a Tie Breaker (TB), you must replace that role with another server before removing the faulty server. The easiest way to do so is to assign that role to a standby node with the same role. A standby MDM can replace an MDM and a standby TB can replace a TB. If there is no other server that can assume these roles, you will need to add a new server, assign that role to it, and only then remove the faulty server.

Procedure

1. Log in to the VxFlex OS GUI.
2. In **Backend** > **Storage** view, locate the faulty server in the MDM cluster and make a note of its details:

Figure 35 Faulty server in **Backend** > **Storage** view

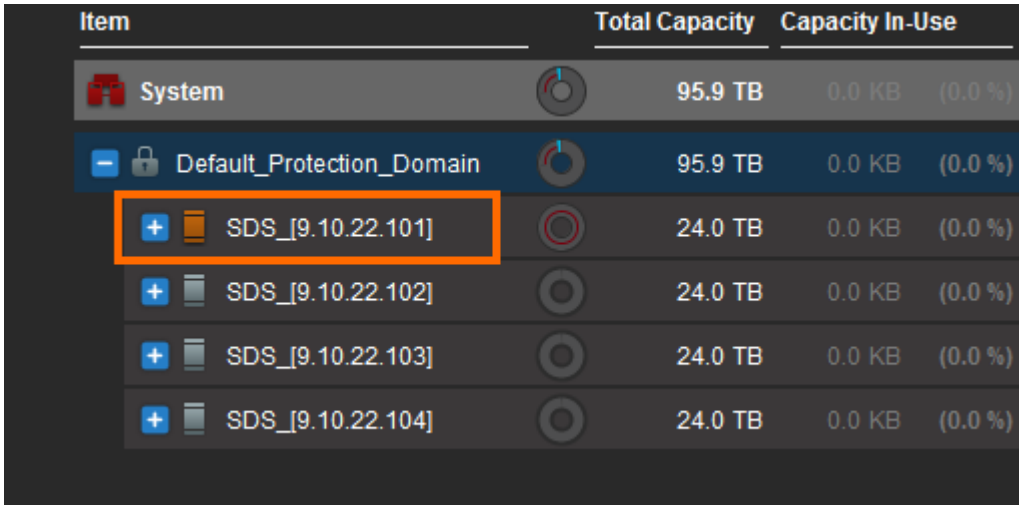

3. In **Hardware** view, in the node list on the left side of the window, click the faulty node that you identified in the previous step.
4. In the node's property sheet on the right side of the window, find the node's management IP address.
5. In the **Node Properties**, if the node has a Standby role, perform the following. Otherwise, continue with step 6 on page 181

   a. Right-click the faulty server and select **Remove Standby**.

      This removes the Standby node from the cluster and removes the MDM.

   b. Verify that all mapped volumes have been removed from the failed node.
6. In **Hardware** view, right-click the faulty server and select **Remove VxFlex OS Software**, then click **OK**.

   This process performs the following:

   - Removes the SDS from the system.
   - Removes the SDS software from the server.
   - Unconfigures and removes the SDC from the server.
   - Reboots the server.
   - Removes the SDC from the system.

   Note

   This may trigger a rebalance. Wait until the rebalance process completes before continuing. This may take over 24 hours, which will cause a timeout error to be displayed for the Remove command.

7. After the rebuild/rebalance completes, repeat the **Remove VxFlex OS Software** procedure in the previous step.
8. When the process completes successfully, a message is displayed. Click **Close**.
9. From the VxFlex OS GUI **Hardware** view, identify the node in the data center:

   a. Right-click the node and select **Set Identification LED**.

b. In the **Set Identification LED** dialog box, select **Turn On**, and then click **OK**.

10. Complete node removal through the AMS.

    a. In the **Nodes** list, right-click the faulty node and select **Remove Node and Unassign IP addresses**.

    This removes the node, and returns its IP addresses to the IP pool so that they can be used by nodes that may join the system later. The node continues to be monitored until it is completely removed.

    b. Click **OK**.

    c. When the process completes successfully, click **Close**.

    The node is removed from the system.

11. For ESXi nodes, from the vSphere client, remove the server from the vCenter.

Adding devices

Add a device to a VxFlex Ready Node system. This procedure is relevant for storage devices and for acceleration devices.

Before you begin

Ensure that an appropriate Storage Pool is defined in the system, because the new device must be assigned to a pool.

If the device that you are adding will be used for DAS Cache, all the SSDs in use with DAS Cache on the specific server must first be removed, and then all SSDs for DAS Cache should be added back to the system together. Up to two DAS Cache devices can be configured on each server.

Procedure

1. If you have not already done so, physically install the device in the chassis.
2. Ensure that the new device was installed correctly and is in the proper state:

    a. From the VxFlex OS GUI **Hardware** view, click the new device.

    b. Find the device's state in the **Physical Drive Properties** sheet.

3. From the **Hardware** view, right-click the device, and choose the required option:

    - For a storage device, choose **Add to Storage Pool**. In the **Add Devices** window, select a Storage Pool. If the device was used previously, select **Force Device Takeover** to overwrite all existing data on the device (all existing data will be lost). Click **OK**.
    - For an acceleration device, choose **Add Acceleration Drive** (this option is only available for SSD and NVMe SSD devices). In the **Add Cache Devices** window, click **OK**

    Note

    In 13G PowerEdge servers, SSDs support both RFcache and DAS Cache. In 14G PowerEdge servers, SSDs support RFcache only. NVMe SSD devices support RFcache only.

    .

Results

When a storage device is added, a rebalance begins. You can monitor progress in the **Dashboard** view. When the rebuild is complete, the new disk has been successfully added to the system.

Removing devices

The removal of storage devices can take a long time, because removal may require data to be moved to other storage devices in the system. If you plan to replace a device with a device containing less storage capacity, you can configure the device to a smaller capacity than its actual capacity, in preparation for replacement. This will reduce rebuilding and rebalancing operations in the system later on. For more information, see "Setting device capacity limits".

The system has job queues for operations that take a long time to execute. You can view the jobs in the **Backend Planned Rebuilds** and **Planned Rebalancing** table views.

There are two ways of removing devices from a system:

- From the **Hardware** view—for details, see "Using the Hardware view to remove and replace a device"
- From the **Backend** view—for details, see "Using the Backend view Remove command". This procedure can only be used to remove a device providing storage. It cannot be used for a device providing caching.

Using the Hardware view to remove and replace a device

This topic explains how to remove a device using the **Remove** command in the Hardware view.

NOTICE

The **Remove** command deletes the specified objects from the system. Use the **Remove** command with caution.

In general, if you want to remove and replace a device, follow these steps:

1. Gather information about the faulty device and the SDS.
2. Remove the faulty device from the VxFlex OS system.
3. Remove the faulty, physical device from the server.
4. Add the replacement device to the server.
5. Add the replacement device to the VxFlex OS system.
6. Ensure that there are no errors.

Identifying a faulty device

The steps to perform are different, depending on whether the device is providing caching (as in the case of some SSD disks), or whether it is used for storage (all HDD disks and some SSD or NVMe SSD disks). Our first step is to locate the faulty device in the VxFlex OS GUI, and then to determine its use.

Procedure

1. Identify the faulty disk:
   a. Log in the VxFlex OS GUI, and from the **Monitor > Alerts** view, find the node that has a faulty disk alert message.

b. From that node's **Hardware** view, write down the following:

- Node information
- Disk WWN
- Disk slot number

c. From the **Hardware** view, locate the node and the faulty disk.

The faulty disk is colored orange.

d. Click the faulty disk, and from the **Physical Drive Properties** section, validate the following:

- The faulty disk's state can be either `Failed`, if the RAID controller has failed it, or `Online (Unused)`, if VxFlex OS has failed it because of too many I/O failures.
- The slot number matches the number previously identified.

2. If the disk is an SSD or NVMe SSD, determine its use:

a. Find the disk's designation in the **Cache** row, underneath the disk:

- RFC - used for RFcache caching
- DC - used for DAS Cache caching (SSD devices only, in 13G PowerEdge servers)
- Blank - used for fast storage (or, if it is orange, may be unused)

The disk usage is now determined: storage or caching.

Results

For HDD, and SSD used for storage, proceed to Removing a storage device on page 185. For SSD used for caching, proceed to Removing a caching device on page 184.

Removing a caching device

This procedure describes how to remove an SSD or NVMe SSD device that is being used to provide caching for acceleration. The caching mode can be Read Flash Cache or DAS Cache (SSD devices, 13G servers only).

Procedure

1. Remove the faulty disk from the VxFlex Ready Node system:

a. From the VxFlex OS GUI **Hardware** view, right-click the faulty disk and choose **Remove Acceleration Drive**.

Note

If you remove a DAS Cache caching device and there is another caching device in the same node, both of them should appear in the **Remove Cache Device** dialog box. All DAS Cache devices displayed will be removed. Later, when adding a replacement SSD device, add all DAS Cache devices together (up to two DAS Cache devices are supported).

---

b. In the **Remove Cache Device** dialog box, click **OK**.

This may result in a rebuild/rebalance cycle. Use the VxFlex OS GUI **Dashboard** view to ensure that the cycle is completed before continuing.

The disk is removed from the VxFlex Ready Node system.

2. Remove the HDD devices that were cached by the faulty device:
   a. From the **Backend > Storage** view, select all the HDDs that show errors, right-click them, and choose **Remove**.

      This removes the disks from the Storage Pool.

   b. In the dialog box, click **OK**.
   c. If a rebuild or rebalance begins, wait for it to finish.
   d. From the **Hardware** view, select all the disks whose state is `Unconfigured Good`, right-click, and choose **Add Storage Device**.

      This adds the disks back to the Storage Pool. The disk state should change to `Online`.

   e. A rebalance begins. Wait for it to finish.

   The disk is removed from the VxFlex Ready Node system.

3. Locate the physical disk:
   a. From the VxFlex OS GUI **Hardware** view, right-click the node, and then select **Set Identification LED**.
   b. In the **Set Identification LED** dialog, select **Turn On**, and then click **OK**.
   c. From the VxFlex OS GUI **Hardware** view, right-click the faulty disk, and repeat the previous step.

      The LEDs flash until turned off.

4. Use the blinking LEDs on the node and disk to identify them in the Data Center.

   Ensure that the node S/N and the disk slot match the information gathered previously. After the disk is removed, stop the blinking LED on the node.

Results

To add a replacement acceleration device, proceed to Adding devices on page 182.

Removing a storage device

This section describes how to remove a device that is being used for storage.

Procedure

1. Remove the faulty disk from the VxFlex Ready Node system:
   a. From the VxFlex OS GUI **Hardware** view, right-click the faulty disk and choose **Remove from Storage Pool**.
   b. In the **Remove Device** dialog, for NVMe SSD devices that you want to physically retain in the chassis, clear the **Prepare Devices for Removal** option. For physical NVMe SSD device removal, leave this option selected.
   c. In the **Remove Device** dialog, for NVMe SSD devices, it is recommended to leave the **Approve reconfigure SVM** option selected.

      A short window of Maintenance Mode may be initiated on the node during the rebuild. If the node contains an MDM, it is preferable that the system is running a 5-node MDM cluster.

   d. For all device types, click **OK**.

When the removal operation is complete, a confirmation message is displayed. If this operation generated a rebuild, it may remain in `Pending` state for an extended period of time.

e. For NVMe SSD devices, if you want to use it for caching or to remove it physically from the chassis, right-click the device and select **Remove Disk Stale Configuration**.

The disk is removed from the VxFlex Ready Node system.

2. Locate the physical disk:

   a. From the VxFlex OS GUI **Hardware** view, right-click the node, and select **Set Identification LED**.

   b. In the **Set Identification LED** dialog, select **Turn On**, and then click **OK**.

   c. From the VxFlex OS GUI **Hardware** view, right-click the faulty disk, and repeat the same steps.

   The LEDs flash until turned off.

3. Use the blinking LEDs on the node and disk to identify them in the Data Center.

   Ensure that the node S/N and the disk slot match the information gathered previously. After the disk is removed, stop the blinking LED on the node.

Results

To add a replacement device, proceed to Adding devices on page 182.

Using the Backend view Remove command

This topic explains how to remove a device using the **Remove** command in the Backend view. This command completely removes the device from the system. This procedure can only be used to remove a device providing storage. It cannot be used for a device providing caching.

Removal operations that are waiting in the job queue are shown as Pending. If a job in the queue is going to take a long time, and you do not want to wait, you can cancel the operation using the **Abort** button in the **Remove** command window (if you left it open), or using the **Abort** command from the **Command menu**.

NOTICE

The **Remove** command deletes the specified objects from the system. Use the **Remove** command with caution.

Procedure

1. In the **Backend** view, navigate to the desired device in the table, and select its row.
2. Right-click the desired device and select the desired **Remove** command.

   a. In the **Remove Device** dialog, for NVMe SSD devices that you want to physically retain in the chassis, clear the **Prepare Devices for Removal** option. For physical NVMe SSD device removal, leave this option selected.

   b. In the **Remove Device** dialog, for NVMe SSD devices, it is recommended to leave the **Approve reconfigure SVM** option selected.

   A short window of Maintenance Mode may be initiated on the node during the rebuild. If the node contains an MDM, it is preferable that the system is running a 5-node MDM cluster.

c. For all device types, click **OK**.

In the confirmation window, click **OK**.The progress of the operation is displayed at the bottom of the window. It is recommended to keep the window open until the operation is completed, and until you can see the result of the operation. For some objects, an **Abort** button is available in the window, which can be used if you decide to abort the operation. There is also an **Abort** command accessible from the **Command menu**.

When the removal operation is complete, a confirmation message is displayed. If this operation generated a rebuild, it may remain in `Pending` state for an extended period of time.

d.

3. Click **Close**.
4. For NVMe SSD devices, if you want to use it for caching or to remove it physically from the chassis, go to the **Hardware** view, right-click the device and select **Remove Disk Stale Configuration**.

   The device is removed from the VxFlex OS system. You may now physically remove the device from the chassis.

5. Locate and remove the physical drive:

   a. From the **Hardware** view, right-click the node, and select **Set Identification LED**.

   b. In the **Set Identification LED** window, select **Turn On**, then click **OK**.

   c. From the **Hardware** view, right-click the faulty device, and repeat the same steps.

   The LEDs flash until they are turned off. The faulty drive can now be physically removed. It is recommended to turn off the drive's LED before removing it from the server.

Turning on identification LEDs

You can locate a specific node or drive in a rack by activating their flashing LEDs from the VxFlex OS GUI. This can help you to identify a specific rack or node for maintenance activities.

Procedure

1. From the **Hardware** view, right-click the node, and select **Set Identification LED**.
2. In the **Set Identification LED** window, select **Turn On**, then click **OK**.
3. From the **Hardware** view, right-click the faulty device, and repeat the same steps.

   The LEDs flash until they are turned off.It is recommended to turn off the flashing LEDs before removing the drive or node.

Adding, removing, and activating and inactivating capacity

Extra storage capacity can be added to VxFlex OS by adding SDSs and/or their storage devices. You can either add them to existing Protection Domains and Storage Pools, or create new ones.

The Dashboard **Capacity** tile, some Backend table views (such as **Capacity Usage**, **Configuration**), and Property Sheets help you to better understand the amount of raw capacity and net free capacity currently available in the system.

Adding, removing, activating, and inactivating Protection Domains

This section describes how to add, remove, activate, and inactivate Protection Domains. Inactivating a Protection Domain does not remove it from the system, but it makes all data stored in that Protection Domain inaccessible to the system.

The inactivation feature is a much more effective way to shut down nodes, and is preferable to shutting them down manually.

When this feature is in effect, the following activities can take place, behind the scenes:

- Determine if there are any current rebuild/rebalance activities taking place. If so, the shutdown will be delayed (unless it is forced) until they are finished.
- Block future rebuild/rebalance activities.
- Quiesce (temporarily disable) application I/O and disable access to volumes.
- Move the DRL mode of all SDSs to harden, in preparation for restarting the server.
- Reload of all SDSs before re-enabling data access.

For each of the following procedures, after you click **OK**, the progress and result of the operation is displayed at the bottom of the window.

Note

Protection Domain removal is only possible if the Protection Domain is empty. If you inactivate a Protection Domain, the data remains on the SDSs, and therefore, it is preferable to remove a Protection Domain if you no longer need it.

Procedure

1. To add a Protection Domain, perform these steps: In the **Backend** > **Storage** view, select the **System** row.
   a. Right-click the row and select **Add Protection Domain**.

      The **Add Protection Domain** window is displayed.

   b. Type a name in the **Name** box, and click **OK**.

      When the operation is complete, the Protection Domain is active. You can now add SDSs, Fault Sets, Storage Pools, and Acceleration Pools to the Protection Domain. Before you add devices, ensure that at least one suitable Storage Pool is defined in the Protection Domain.

2. To remove Protection Domains, perform these steps:
   a. In the **Backend** > **Storage** view, navigate to, and select one or more Protection Domains.
   b. Verify that you have removed all child nodes from the Protection Domain.
   c. Right-click the Protection Domain and select **Remove**.
   d. Click **OK**.
   e. If a confirmation window appears, confirm the operation, and type your password if requested to do so.

3. To inactivate Protection Domains, perform these steps:
   a. In the **Backend** > **Storage** view, navigate to, and select one or more Protection Domains.
   b. Right-click the Protection Domain and select **Inactivate Protection Domain**.

      The **Inactivate Protection Domain** window is displayed.
   c. Click **OK**.

      If a confirmation window appears, confirm the operation, and type your password if requested to do so.
4. To activate Protection Domains, perform these steps:
   a. In the **Backend** > **Storage** view, navigate to, and select one or more Protection Domains.
   b. Right-click the Protection Domain and select **Activate Protection Domain**.

      The **Activate Protection Domain** window is displayed.
   c. Click **OK**.

Adding Fault Sets

Fault Sets provide additional safeguards for protecting your data against hardware failure. Fault Sets are subsets of a given Protection Domain.

Note

When defining Fault Sets, you must follow the guidelines described in Fault Sets on page 33. Failure to do so may prevent creation of volumes.

Procedure

1. In the **Backend** > **Storage** view, navigate to, and select the Protection Domain.
2. Right-click the Protection Domain and select **Add Fault Set**.
3. Type a name in the **Fault Set Name** box, and click **OK**.

   The Fault Set will now be visible in the **Related Objects** section of the Protection Domain's Property Sheet.

   Note

   Use the CLI to remove Fault Sets.

Adding Storage Pools

A Storage Pool is a group of devices within a Protection Domain. Create Storage Pools before you start adding devices to the system. Each time that you add devices to the system, you must map them to either Storage Pools.

Note

You cannot enable zero padding after adding the devices. For more information, see Storage Pools on page 31.

Procedure

1. In the **Backend** view, select the desired Protection Domain.
2. From the **Command menu** or context-sensitive menu, select **Add Storage Pool**.

   The **Add Storage Pool** window is displayed.
3. Type a name in the **Name** box.

   Note

   The RAM Read Cache features are advanced features, and it is usually recommended to accept the default values. You can configure these features later, if necessary, using the `Configure RAM Read Cache` command. For more information about RAM Read Cache features, see Managing read RAM cache on page 116
4. Click **OK**.

   To remove a Storage Pool, perform these steps:
5. In the **Backend** view, select the Storage Pool.
6. From the **Command menu** or context-sensitive menu, select **Remove**.

   The **Remove Storage Pool** window is displayed.
7. Click **OK**.

Activating devices

Use the `Activate Device` command in the following situations:

- Storage devices were added to the system using the **Test only** option for Device Tests, and successfully passed the tests.
- Storage devices were inactivated, and you want to bring them back online.

Procedure

1. In the **Backend** > **Storage** view, navigate to the device or devices in the table, and select the corresponding rows.
2. Right-click and select **Activate Device**.

Clearing device errors

Procedure

1. In the **Backend** > **Storage** view, navigate to the device in the table, and select its row.
2. Right-click the row and select **Clear Device Errors**.

Setting device capacity limits

In circumstances where you need to replace a storage device in your system with a storage device of a smaller capacity, you should first set the capacity limit of the device to be removed to less than its full capacity. In such a case, capacity will be decreased, but the size of the disk remains unchanged. The capacity assigned to the storage device must be smaller than its actual physical size.

Note

Decreased capacity is shown on the **Dashboard**, using pale gray, on the outer ring on the Capacity tile.

Procedure

1. In the **Backend** > **Storage**view, navigate to the device in the table, and select its row.
2. Right-click the row and select **Set Device Capacity Limit**.
3. Type the desired value and click **OK**.

Modifying cache settings

The following SSD caching modes are available:

- Read Flash Cache
- DAS Cache

The considerations for using each mode are explained in Caching on page 38. Enabling caching is a two-step process: enable caching at the Storage Pool level, and then designate individual SSDs as caching devices. Within a Storage Pool, only one mode of caching can be enabled.

A maximum of two SSDs can be configured as DAS Cache.

Note

Before performing any cache removal or addition operations, it is recommended to stop caching on the Storage Pool, and start caching again after those operations are complete.

This topic includes the following procedures:

- "Enabling and disabling caching at Storage Pool level"
- "Designating an SSD as a caching device"
- "Stopping caching on a specific device"
- "Removing an SSD device from the caching pool"
- "Configuring Read Flash Cache settings per node"

Enabling and disabling caching at Storage Pool level

To enable or disable caching in a Storage Pool, follow these steps:

Procedure

1. In the **Backend** view, locate and right-click the required Storage Pool, and choose either **Set Read Flash Cache Caching Policy** or **Set DAS Cache Caching Policy**, depending on your server model, and the type of caching that you wish to use.
2. In the window that appears:
   - To enable caching, select the **Use Read Flash Cache/DAS Cache** option, and click **OK**.

- To disable caching, clear the **Use Read Flash Cache/DAS Cache** option, and click **OK**.

Designating an SSD as a caching device

Note

A maximum of two SSDs can be configured as DAS Cache.

To designate an SSD as a caching device, follow these steps:

Procedure

1. If the device is currently being used for storage, remove the device from the Storage Pool, as described in "Procedure—Removing a storage device".
2. Ensure that caching is enabled at the Storage Pool level, as described in "Enabling and disabling caching at Storage Pool level".
3. In the **Hardware** view, right-click the SSD device required for caching, and choose **Add Cache Drive**.
4. In the **Add Cache Devices** window, select the **Cache** type that you enabled for the Storage Pool, and click **OK**.

   The SSD device is now configured as a caching device, and will no longer appear in the list of devices shown in the **Backend** views. In the **Hardware** view, the type of caching provided by the device is shown below each device, and in the associated **Physical Device Properties** sheet.

Stopping caching on a specific device

To stop caching on a specific SSD device:

Procedure

1. In the **Hardware** view, right-click the SSD device, and choose **Remove Cache Drive**.
2. In the **Remove Cache Devices** window, click **OK**.

   The SSD device has been removed from the caching pool. To use the SSD for storage, follow the instructions in "Adding devices".

Removing an SSD device from the caching pool

This procedure describes how to remove an SSD device that is being used to provide caching. You may want to do this in order to use the SSD for storage, instead of for caching. If you want to completely remove the SSD from the system, follow the instructions in "Using the Hardware view to remove and replace a device".
The caching mode can be Read Flash Cache, DAS Cache, or CacheCade. In this procedure, where the instructions apply equally to all forms of caching, we will use the phrase <Caching_Method>, that should be replaced with either "Read Flash Cache", "DAS Cade", or "CacheCade", corresponding to your system configuration.

Procedure

1. From the **Backend** view, right-click the Storage Pool of which the disk is a member, and choose **Set <Caching_Method> Caching Policy**.

   It is easiest to do this by sorting the SDSs by Storage Policy.

2. In the **Set <Caching_Method> Caching Policy** window, select **Don't use <Caching_Method>**, then click **OK**.
3. When the status shows `Does not use cache`, click **Close**.
4. From the **Hardware** view, right-click the faulty disk and choose **Remove Cache Drive**.
5. In the **Remove Device** window, click **OK**.
6. In some cases, a rebuild/rebalance begins.

   Use the **Dashboard** view to determine when that is complete.
7. Clear disk errors.
8. To add the SSD to a Storage Pool, refer to "Adding devices".

Configuring Read Flash Cache settings per Protection Domain

This procedure explains how to configure some Read Flash Cache-specific settings. These settings are applied to Read Flash Cache devices installed in a specific Protection Domain. Settings include:

- **Page Size** is the page size used for caching on the flash device, in multiples of 4 KB (4 KB–64 KB).
- **Max I/O Size** is the maximum I/O size in KB.
- **Pass-Through Mode** is disabled by default, and should only be used in certain use cases, because it affects performance. Options include: **Invalid**, **None**, **Read**, **Write**, **Read and Write**, or **Write Miss**

Procedure

1. From the **Backend** view, right-click the Protection Domain, and select **Configure Read Flash Cache**.
2. In the **Configure Read Flash Cache** window, you can do the following:
   a. To change the page size used by Read Flash Cache, select the required value from the drop-down list.
   b. To set maximum I/O size, edit the contents of the **Max I/O Size** box.
   c. To set the pass-through mode, select one of the options from the **Pass-Through Mode** drop-down list.
3. Click **OK**.

Changing Read RAM Cache volume settings

By default, Read RAM Cache is disabled on volumes. To change Read RAM Cache settings on volumes, perform these steps:

Procedure

1. In **Frontend** > **Volumes**, select the **Volumes** or **Volumes Monitor** view.
2. Navigate to the volumes, and select them.
3. Right-click the volumes and select **Set Volume Read RAM Cache**.

   The **Set Volume Use Read RAM Cache** window is displayed, showing a list of the volumes that will be modified.
4. Select or clear the **Use Read RAM Cache** check box as follows:

- To disable Read RAM Cache on the volumes, clear the check box.
- To enable Read RAM Cache on the volumes, select the check box.

5. Click **OK**.

   The progress of the operation is displayed at the bottom of the window. It is recommended to keep the window open until the operation is completed, and until you can see the result of the operation.

   Figure 36 Set Volume Use Read RAM Cache window

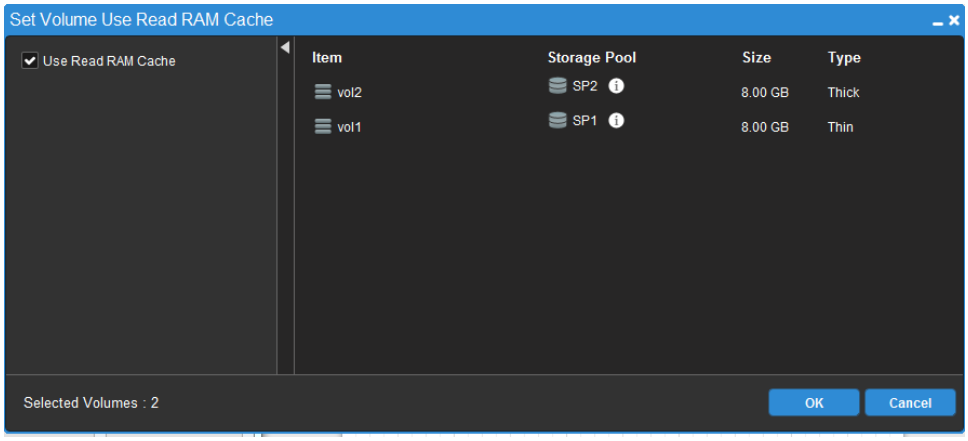

Clearing device errors

Procedure

1. In the **Backend** > **Storage** view, navigate to the device in the table, and select its row.
2. Right-click the row and select **Clear Device Errors**.

Configuring Read RAM Cache (advanced, Backend)

The Read RAM Cache feature improves your system's application performance for storage-related activities. By default, caching is disabled.

To use Read RAM Cache, you need to configure Read RAM Cache settings at two levels:

- Storage Pool—controls Read RAM Cache for all the SDSs in the selected Storage Pool. Caching can be enabled or disabled, and either Cached (default) or Passthrough Write Handling modes can be selected. When Read RAM Cache is enabled in a Storage Pool, the feature is enabled at Storage Pool level. However, caching must also be set to Enabled in each SDS in the Storage Pool. Caching will only begin once devices have been added to the SDSs. It is possible to enable RAM caching for a Storage Pool and then disable caching on one or more SDSs individually.
- Per SDS—controls Read RAM Cache for one or more SDSs. Caching can be enabled or disabled for the specified SDS, and the capacity allocated for caching on an SDS can be specified. Caching will only begin after one or more devices are added to the SDSs. Ensure that the feature is also enabled at Storage Pool level.

  To configure caching, perform these steps:

Procedure

1. In the **Backend** > **Storage** view, navigate to, and select the desired Storage Pools.
2. From the **Command menu** or context-sensitive menu, select **Configure Read RAM Cache**.

   The **Configure Read RAM Cache** window is displayed. The right pane of the window lists the Storage Pools that you are configuring.
3. Select or clear the options that you require (selected=used; clear=not used), and click **OK**.

   Figure 37 Read RAM Cache configuration at Storage Pool level

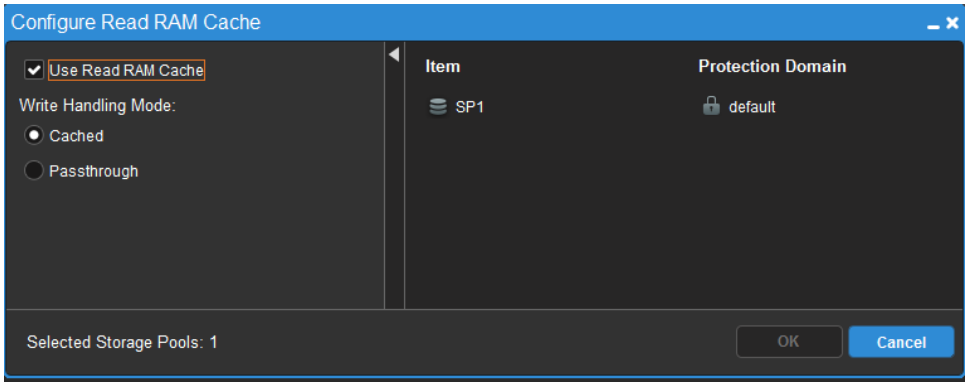

4. To enable\disable\configure cache size for SDSs, in the **Backend** view, navigate to, and select the desired SDS(s).
5. From the **Command menu** or context-sensitive menu, select **Configure Read RAM Cache**.

   The **Configure RAM Read Cache** window is displayed. The right pane of the window lists the SDSs that you are configuring.
6. Select or clear the option that you require (selected=enable; clear=disable).
7. If necessary, edit the value in the **Read RAM Cache Size** box (default=128 MB).
8. Click **OK**.

   Figure 38 Read RAM Cache configuration at SDS level

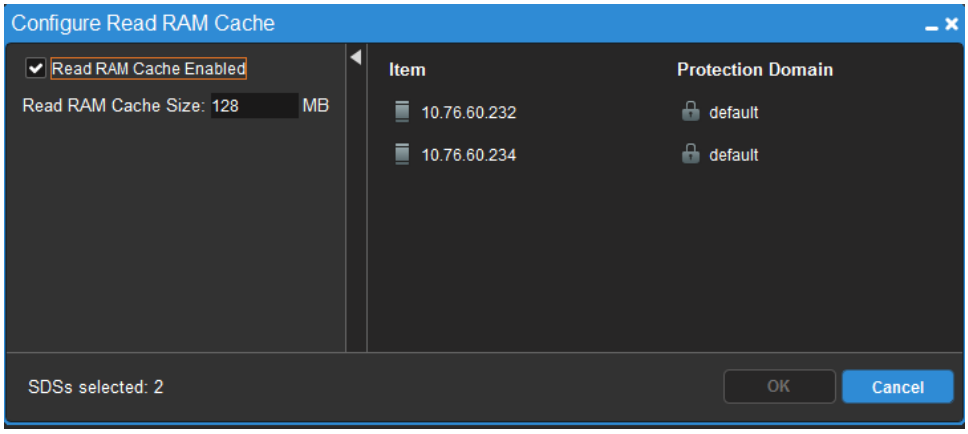

Configuring volumes, volume trees, SDCs, and snapshots

This section contains procedures for adding, removing, and managing volumes and snapshots. It explains how to remove volumes, create snapshots, and set volume bandwidth and IOPS limits. It also describes setting the SDC restriction mode and how to approve SDCs before mapping volumes.

Adding volumes

Add volumes to a system.

Before you begin

There must be at least three SDS nodes in the system and there must be sufficient capacity available.

Note

For the minimum size of an SDS, see System requirements on page 21.

The adding and mapping volume process is necessary, as part of the getting started process, before applications can access the volumes. In addition, you may create additional volumes and map them as part of the maintenance of the virtualization layer.

You can configure the caching option when creating the volumes, or you can change the Read RAM Cache feature later. If you want to enable the caching feature, ensure that the feature is also enabled in the backend of the system, for the corresponding Storage Pool and SDSs. For more information, see Changing Read RAM Cache volume settings on page 193.

Define volume names according to the following rules:

- Contains less than 32 characters
- Contains only alphanumeric and punctuation characters
- Is unique within the object type

VxFlex OS objects are assigned a unique ID that can be used to identify the object in CLI commands. You can retrieve the ID via a query, or via the object's property sheet in the VxFlex OS GUI. It is highly recommended to give each volume a meaningful name associated with its operational role.

To add one or multiple volumes, perform these steps:

Procedure

1. In any of the **Frontend** > **Volumes** views, navigate to the Storage Pool to which you want to add the volume, and select it.
2. From the **Command menu** or context-sensitive menu, select **Add Volume**.
3. In the **Add Volume** window, if you want to create more than one volume, select **Create multiple volumes** and type the number of volumes you would like to add in the **Copies** box.
   - If you type `1`, only one volume will be created (optional—can be left blank).
   - If you type a number greater than 1, the characters `%i%` will be added to the **Name** box, and multiple volumes will be created, accordingly.
     The volumes will be named and numbered automatically, starting from 1. If you want the numbering to start from a different number, type it in the

Start numbering at box, as described in Step 5. The remaining options in the window will be assigned to all the volumes created in this operation.

4. Type a name for the volume:
   - If you are adding one volume, enter the name in the **Name** box.
   - If you are adding multiple volumes, enter the base name in the **Base name** box.
     The volumes will all be created with the same name, and a number will be appended instead of the characters `%i%`. These characters can be positioned anywhere in the name. The names that will be created are displayed in the right pane of the window, as shown in the figure later in this topic.
5. If you want the numbering to start from a specific number other than 1, type it in the **Start numbering at** box.

   This number will be the first number in the series that will be appended to the volume name. For example, if the **Name** is `Vol%i%` and the **Start numbering at** value is `100`, the name of the first volume created will be Vol100, and the second volume will be Vol101, and so on.
6. Type a number in the **Size** box, representing the volume size in GB (basic allocation granularity is 8 GB).
7. Select either **Thick** (default) or **Thin** provisioning options.
8. If you want to enable the RMcache feature (disabled by default), select **Use RMcache**.
9. Click **OK**.

   The progress of the operation is displayed at the bottom of the window. It is recommended to keep the window open until the operation is completed, and until you can see the result of the operation.

Figure 39 Add Volume window

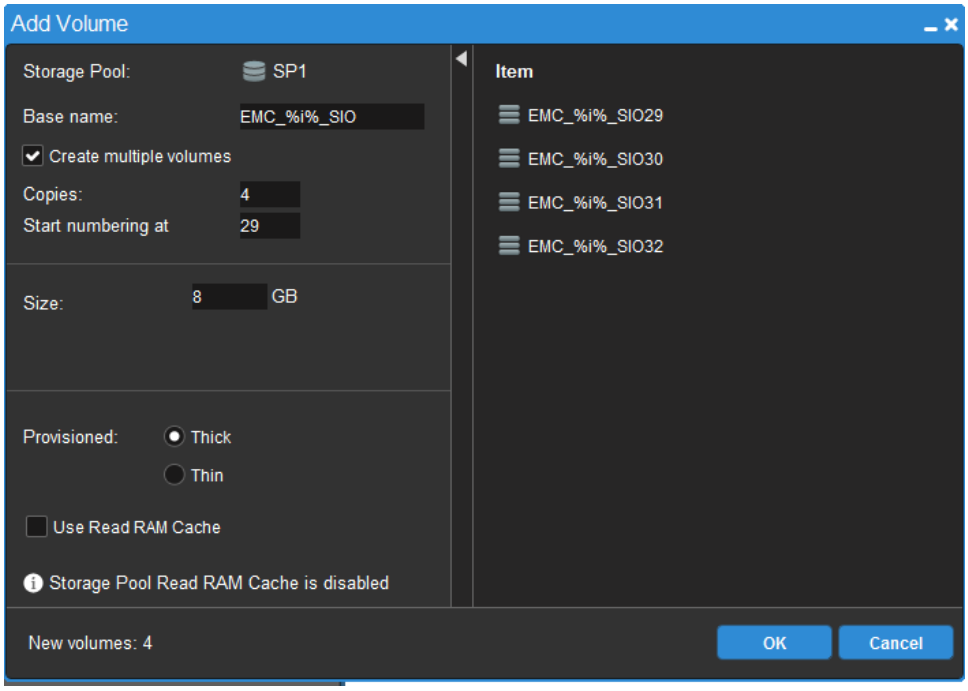

After you finish

To use the created volume, you must map it to (at least) one SDC. If the restricted SDC mode is enabled for the system, you must approve SDCs prior to mapping volumes to them. For more information on approving SDCs, see Approve SDCs using the VxFlex OS GUI on page 199. For more information on mapping volumes, see Mapping a volume to an SDC on page 94.

Restricted SDC mode

Enabling restricted SDC mode gives you more control over VxFlex OS.

Restricted SDC mode is set at the system level. When enabled, you must approve SDCs prior to mapping volumes to them. The restricted SDC setting has the following modes:

- No restriction — Volumes can be mapped without pre-approving SDCs.
- GUID restriction — SDCs are approved using their GUID. Only once the SDC is approved can it be used for mapping volumes.
- Approved IP restriction — SDCs must be approved using their GUID and IP address. Only once the SDC is approved using both GUID and IP address can it be used for mapping volumes.

You can set the restricted SDC mode using the VxFlex OS GUI, CLI, or REST API.

Note

In a system that has been upgraded and already has volumes mapped to SDCs, if you want to enable restricted SDC mode, you must first approve the SDCs and only then enable restricted SDC mode.

Set the system's restricted SDC mode using the VxFlex OS GUI

Use the VxFlex OS GUI to set the restricted SDC mode.

The system's restricted SDC mode can also be set using the CLI or the REST API. For details, see the *CLI Reference Guide* or the *REST API Reference Guide*.

Note

In a system that has been upgraded and already has volumes mapped to SDCs, if you want to enable restricted SDC mode, you must first approve the SDCs and only then enable restricted SDC mode.

Procedure

1. In the **Frontend** > **SDCs** view, right-click on the System and select **Set Restricted SDC Mode**.
2. In the **Set Restricted SDC Mode** dialog box, select one of the following options and click **OK**.
   - No restriction
   - GUID restriction
   - Approved IP restriction

Results

The restricted SDC mode is set for the system. If you enabled restricted SDC mode by selecting either `GUID restriction` or `Approved IP restriction`, you must configure approved SDCs before you can map volumes.

Approve SDCs using the VxFlex OS GUI

When the system's restricted SDC mode is set to GUID restriction, you must approve SDCs before you can map volumes.

Procedure

1. In the **Frontend** > **SDCs** view, right-click on one or several SDCs that you want to approve and select **Approve SDC**.
2. In the **Approve SDC** window, verify that the SDCs listed are the ones you want to approve and click **OK**.

Results

The SDCs are approved and you can map volumes. The **Approved IPs** column in the **Frontend** > **SDCs** displays which SDCs are approved.

Configure approved SDC IP addresses using the VxFlex OS GUI

When the system's restricted SDC mode is set to `approved IP restriction`, you must configure SDC IP addresses before you can map volumes.

Before you begin

Ensure that the SDCs have been approved by GUID.

Procedure

1. In the **Frontend** > **SDCs** view, right-click on the SDCs that you want to approve and select **Configure Approved IP Addresses**.
2. In the **Configure Approved IP Addresses** window, add the IP addresses of the SDCs you want to approve.

   You can click the **Add IP Address** button to add up to a total of four IP addresses.
3. Click **OK** and then click **Close**.

Results

The SDC IP addresses are approved and you can map volumes. The **Approved IPs** column in the **Frontend** > **SDCs** view displays which SDC are approved.

Mapping volumes

Map one or more volumes to/from SDCs. Mapping exposes the volume to the specified SDC, effectively creating a block device on the SDC.

For Linux devices, the `scini` device name can change on reboot. It is recommended to mount a mapped volume to the VxFlex OS unique ID, a persistent device name, rather than to the `scini` device name.

To identify the unique ID, run the `ls -l /dev/disk/by-id/` command. For more information, see Associating VxFlex OS volumes with physical disks on page 248. You can also identify the unique ID using VMware. In the VMware management interface, device is called **EMC Fibre Channel Disk**, followed by an ID number starting with the prefix **eui**.

To map volumes, perform these steps:

Procedure

1. In any of the **Frontend** > **Volumes** views, navigate to the volumes, and select them.
2. From the **Command menu** or context-sensitive menu, select **Map Volumes**.

   The **Map Volumes** window is displayed, showing a list of the volumes that will be mapped.
3. In the **Select Nodes** panel, select one or more SDCs to which you want to map the volumes.
   - You can use the search box to find SDCs.
   - If you select an SDC that is already mapped to the volume, a green icon will appear in the mapping matrix on the right side of the window.
4. Click **Map Volumes**.

   The progress of the operation is displayed at the bottom of the window. It is recommended to keep the window open until the operation is completed, and until you can see the result of the operation.

Unmapping volumes

Unmap one or more volumes from SDCs.

Procedure

1. In any of the **Frontend** > **Volume** views, navigate to the volumes, and select them.
2. From the **Command menu** or context-sensitive menu, select **Unmap Volumes**.

   The **Unmap Volumes** window is displayed, showing a list of the volumes that will be unmapped.

   On ESX servers only, by default, the system checks that the volume is not in use on any SDC/ESX. The **Force** option will make the command ignore this usage check, and the operation therefore work faster.
3. If you want to exclude some SDCs from the unmap operation, in the **Select Nodes** panel, select one or more SDCs for which you want to retain mapping.

   You can use the search box to find SDCs
4. Click **Unmap Volumes**.

   The progress of the operation is displayed at the bottom of the window. It is recommended to keep the window open until the operation is completed, and until you can see the result of the operation.

Removing volumes

Before you begin

Ensure that the volume you are removing is not mapped to any SDCs. If it is, unmap it before removing it. For information, see Unmapping volumes on page 200.

If you want to remove a volume's related snapshots, or just the snapshots, see Removing snapshots on page 203.

Best practice is to avoid deleting volumes or snapshots while the MDM cluster is being upgraded, to avoid causing a Data Unavailability status.

Note

Removal of a volume erases all the data on the corresponding volume.

Procedure

1. In any of the views of the **Frontend** > **Volumes** view, navigate to the volume or volumes you want to remove, and select them.
2. From the **Command menu** or context-sensitive menu, select **Remove**.

   The **Remove Volumes** window is displayed, showing a list of the volumes that will be removed.
3. Click **OK**.

   The progress of the operation is displayed at the bottom of the window. It is recommended to keep the window open until the operation is completed, and until you can see the result of the operation.

Creating a datastore

When there are SDC components on ESXi servers, in order to be able to use the storage on the ESX, you must create a datastore to store the virtual machines that are running on the host.

Before you begin

- Ensure that the vCenter is configured in the AMS system settings. The vCenter can be configured during AMS deployment, or at any time in the **System Settings** menu.
- Verify that all hosts are connected to the vCenter.

Note

It is critical that all hosts be connected to the vCenter.

Procedure

1. In the **Frontend** > **Volumes** view, right-click on a Storage Pool and select **Create Datastores**.

   The **Datastores Creation** window is displayed.
2. In the **Create Volumes** step, if you want to create more than one volume, select **Create multiple volumes**.

   a. Type the number of volumes you would like to add in the **Copies** box.

   b. If you type `1`, only one volume will be created (optional—can be left blank).

   c. If you type a number greater than 1, the characters `%i%` will be added to the **Name** box, and multiple volumes will be created, accordingly.

      The volumes will be named and numbered automatically, starting from 1. The remaining options in the window will be assigned to all the volumes created in this operation.
3. Type a name in the **Name** box.

   If you are adding multiple volumes, they will all be created with the same name, and a number will be appended instead of the characters `%i%`.

   These characters can be positioned anywhere in the name. The names that will be created are displayed in the right pane of the window.

4. Type a number in the **Size** box, representing the volume size in GB (basic allocation granularity is 8 GB).
5. Select either **Thick** (default) or **Thin** provisioning options.
6. If you want to enable the RAM Read Cache feature, select **Use RAM Read Cache**.
7. Click **Create**.

   The **Map Volumes** step is displayed, showing a list of the volumes that will be mapped.
8. In the **Select Nodes** panel, select one or more SDCs to which you want to map the volumes.
   - You can use the search box to find SDCs.
   - If you select an SDC that is already mapped to the volume, a green icon will appear in the mapping matrix on the right side of the window.
9. Click **Map Volumes**.

   The **Create Datastores** step is displayed.
10. In the **Create Datastores** step, type the datastore name.

    The datastore name is a combination of the entered datastore name, and the characters %vol%. If you are creating multiple datastores, the characters %i% are appended to indicate the index number of the datastore.
11. From the **File System Type** list, select **VMFS**.
12. Select **Create**.

Deleting a datastore

Procedure

1. In the **Frontend** > **Compute** view, select the **Datastore** view.
2. Right-click the datastore you want to delete and select **Delete Datastore**.
3. In the **Delete Datastore** window, verify that the datastore listed is the one you intend to delete.
4. Click **Delete**.

Creating volume snapshots

Take a snapshot of one or more volumes.

This topic describes how to take a snapshot of one or more volumes.

When specifying more than one volume (a list), a consistency group is generated by default, and can be viewed in the snapshot's property sheet. The snapshots under the consistency group are taken simultaneously for all listed volumes, thus ensuring their consistency.

You can accept the default name, or define snapshot names according to the following rules:

- Contains less than 32 characters
- Contains only alphanumeric and punctuation characters
- Is unique within the object type

VxFlex OS objects are assigned a unique ID that can be used to identify the object in CLI commands. You can retrieve the ID via a query, or via the object's property sheet in the VxFlex OS GUI.

Note

The consistency group is for convenience purposes ONLY. There are no protection measures to conserve the consistency group. You can delete members from it.

To take a snapshot, perform these steps:

Procedure

1. In the **Frontend** > **Volumes** view, navigate to the volumes, and select them.
2. From the **Command menu** or context-sensitive menu, select **Snapshot Volume**.

   The **Snapshot Volume** window is displayed, showing the volumes for which snapshots will be created.
3. In the **Index** box, type the number that you want to append to the snapshot names.
4. If you want the snapshots to belong to a consistency group, ensure that the **Create Consistency Group** check box is selected.
5. Click **OK**.

   The progress of the operation is displayed at the bottom of the window. It is recommended to keep the window open until the operation is completed, and until you can see the result of the operation.

Removing snapshots

You can remove a volume together with its snapshots, or remove individual snapshots. Before removing a volume or snapshots, you must ensure that they are not mapped to any SDCs. If they are, unmap them before removing them. Snapshots are unmapped in the same way as volumes are unmapped. For information, see Unmapping volumes on page 200.

Best practice is to avoid deleting volumes or snapshots while the MDM cluster is being upgraded, to avoid causing a Data Unavailability status.

Note

Removal of a volume or snapshot erases all the data on the corresponding volume or snapshot.

Procedure

1. In the **Frontend** > **Volumes** > **V-Trees** view, navigate to the volume from which you want to remove snapshots, and select it.
2. From the **Command menu** or context-sensitive menu, select one of the following options:
   - To remove both the parent volume and all volumes that were created as snapshots of the specified volume or one of its descendants, select **Remove with Descendants**
   - To retain the parent volume, and remove only its snapshots, select **Remove Descendants Only**

The **Remove Volumes** window is displayed, showing a list of the objects that will be removed.

3. Click **OK**.

   The progress of the operation is displayed at the bottom of the window. It is recommended to keep the window open until the operation is completed, and until you can see the result of the operation.

Removing snapshots from a consistency group

Remove a snapshot from a consistency group using the VxFlex OS GUI.

Procedure

1. In the **Frontend** > **Volumes** > **V-Trees** view, navigate to the snapshot which you want to remove from the consistency group, and select the snapshot.
2. From the **Command menu** or context-sensitive menu, select **Remove Consistency Group**.

   The **Remove Consistency Group** window is displayed, showing the selected snapshot.
3. Click **OK**.

   The progress of the operation is displayed at the bottom of the window. It is recommended to keep the window open until the operation is completed, and until you can see the result of the operation.

Increasing a volume's size

You can increase (but not decrease) the capacity of one or more volumes at any time, as long as there is enough capacity for the volume size to grow.

Procedure

1. In **Frontend** > **Volumes**, select the **Volumes** or **Volumes Monitor** view.
2. Navigate to the volumes, and select them.
3. From the **Command menu** or context-sensitive menu, select **Increase Volumes' Size**.

   The **Increase Volumes' Size** window is displayed, showing a list of the volumes that will be modified.
4. In the **New Size** box, type a number representing the new volume size in GB (basic allocation granularity is 8 GB).
5. Click **OK**.

   The progress of the operation is displayed at the bottom of the window. It is recommended to keep the window open until the operation is completed, and until you can see the result of the operation.

Setting volume bandwidth and IOPS limits

Set bandwidth and IOPS limits for volumes on a per-SDC basis. This enables you to control the quality of service (QoS).

Before you begin

Ensure that the volumes are mapped before you set these limits.

Procedure

1. In **Frontend** > **Volumes** view, select the **Volumes** or **Volumes Monitor** view.
2. Navigate to the volumes, and select them.
3. From the **Command menu** or context-sensitive menu, select **Set Volume Limits**.

   The **Set Volume Limits** window is displayed, showing a list of the volumes to be modified.
4. In the **Bandwidth Limits** and **IOPS Limits** boxes, type the required values, or select the corresponding **Unlimited** check box.
   - The number of IOPS must be larger than 10.
   - The volume network bandwidth is in MB/sec.
5. In the **Select Nodes** panel, select the SDCs to which you want to apply the changes.
6. Click **Set Limits**.

   The progress of the operation is displayed at the bottom of the window. It is recommended to keep the window open until the operation is completed, and until you can see the result of the operation.

Entering and exiting SDS Instant Maintenance Mode

This topic explains how to put an SDS into Maintenance Mode, in order to perform non-disruptive maintenance on the SDS, and how to cancel Maintenance Mode when you are finished. Instant Maintenance Mode lets you restart a server that hosts an SDS, without initiating data migration or exposing the system to the danger of having only a single copy of data. The system displays the SDSs that are in Maintenance Mode at any given time (but does not provide the total number of SDSs).

While SDSs are in Maintenance Mode, you should avoid both unnecessary rebuilds and operations that require taking SDS offline temporarily. It is recommended to use Maintenance Mode when there is relatively low system activity, as the time it takes for an SDS to exit Maintenance Mode depends on the amount of data that needs to be synchronized back into the server.

Note

Functional operations, such as configuration, cannot be performed on an SDS while it is in Maintenance Mode. If an active full copy is lost, its data will be unavailable until the SDS is brought back into the system, but that data will not be lost; it will be rebuilt using the temporary copy.

Procedure

1. To put an SDS into Maintenance Mode, perform these steps:
   a. In the **Backend > Storage** view, navigate to, and select the desired SDS.
   b. From the **Command menu** or context-sensitive menu, select **Enter Maintenance Mode**.

      The **Enter Maintenance Mode** window is displayed.

c. If you want to force entry into Maintenance Mode even though there is insufficient space or degraded/failed capacity, select the corresponding check box:

- **Force Insufficient Space**—allow entry into maintenance mode, even without enough available capacity
- **Force Degraded or Failed**—allow entry into maintenance mode, even with degraded or failed data

d. Click **OK**.

The status area at the bottom of the window indicates when the operation is complete. Once the SDS is in Maintenance Mode, this will be indicated both on the Dashboard, and in Backend tables and Property Sheets, using the symbol, and the Maintenance Mode color code (green).

2. To put an SDS back into regular service (cancel Maintenance Mode), perform these steps:

a. In the **Backend > Storage** view, navigate to, and select the desired SDS.

b. From the **Command menu** or context-sensitive menu, select **Exit Maintenance Mode**.

The **Exit Maintenance Mode** window is displayed.

c. If you want to force exit from Maintenance Mode even though there is a failed SDS, select the **Force Failed SDS** check box.

d. Click **OK**.

The status area at the bottom of the window indicates when the operation is complete. Once the operation has been successfully completed, the SDS returns to normal operation, and data deltas collected on other SDSs during the maintenance period are copied back to the SDS.

Cluster operations

Use the VxFlex OS GUI **Hardware** view to perform cluster operations such as adding and removing standby servers, adding and removing Managers and TieBreakers, and replacing a Master MDM.

The following cluster operations can be performed using the **Cluster** menu in the **Hardware** view:

- Adding a standby server, as described in Adding a standby to the MDM cluster on page 207
- Removing a standby server, as described in Removing a standby from the MDM cluster on page 207
- Extending a cluster, as described in Extending an MDM cluster on page 208
- Reducing a 5-node cluster to a 3-node cluster, as described in Switching from 5-node cluster mode to 3-node cluster mode on page 209
- Replacing a Master MDM, as described in Replacing a Master MDM on page 209
- Transfer the MDM cluster component using the AMS on page 110

Adding a standby to the MDM cluster

Standby servers can be configured in the system, in preparation for adding them to an MDM cluster.

Before you begin

Ensure that you know the name of the node that you want to add as a standby to the MDM cluster.

The MDM cluster can only be extended by promoting existing standby servers already defined in the system. Standby servers for the MDM cluster can only be defined after system deployment. To add standby servers to the cluster, perform the following steps.

Procedure

1. Using the **Hardware** view, find a server that does not currently have an MDM role. In the **Property Sheet**, **Node Properties** section, find a node where the VxFlex OS Role is "None MDM".

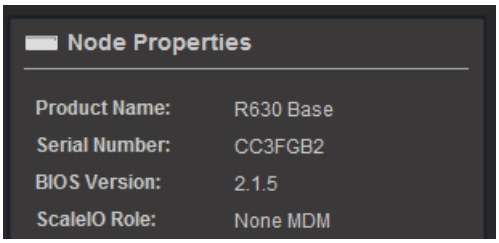

2. In the **Hardware** view, in the **Nodes** list on the left side, right-click the required node, and select **Add as Standby**.

   The **Add Standby Node** window is displayed.

3. In the **Standby Name** box, type a name for the standby.
4. From the **Role** list, select one of the following:

Option	Description
MDM	Creates a standby with the MDM role
TieBreaker	Creates a standby with the Tie Breaker role

5. Click **OK**.

Results

The MDM/Tie Breaker package is installed on the node, and the standby is added to the MDM cluster. To change a standby server's status to either MDM or Tie Breaker, follow the instructions in Extending an MDM cluster on page 208.

Removing a standby from the MDM cluster

A server can be gracefully removed from the MDM cluster if the server is in standby state.

Once a server is in standby state, it can be gracefully removed from the system without impacting on normal system operation.

Procedure

1. In the **Hardware** view, click the **Cluster** menu, and select **Remove Standbys**.

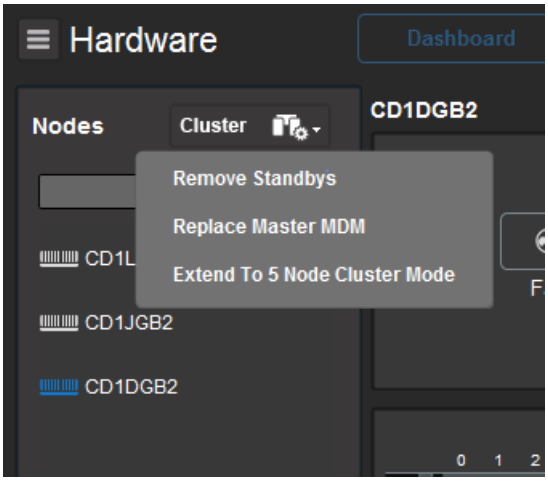

   The **Remove Standby** window is displayed.

2. Select the standbys that you want to remove, and click **OK**.

Results

The selected standby servers are removed from the MDM cluster, and the MDM packages are removed from those servers.

Extending an MDM cluster

A 3-node cluster can be enlarged to a 5-node cluster, to provide more protection for the MDM cluster.

Before you begin

Ensure that a standby MDM and a standby Tie Breaker are already defined in the MDM cluster. For information about adding standby servers, see Adding a standby to the MDM cluster on page 207.

A 5-node cluster provides better high availability than a 3-node cluster, and therefore better protection in fault scenarios. Standby servers can be added to a 3-node cluster at any time, using the VxFlex OS GUI.

Procedure

1. In the **Hardware** view, click the **Cluster** menu, and select **Extend to 5 Node Cluster Mode**.

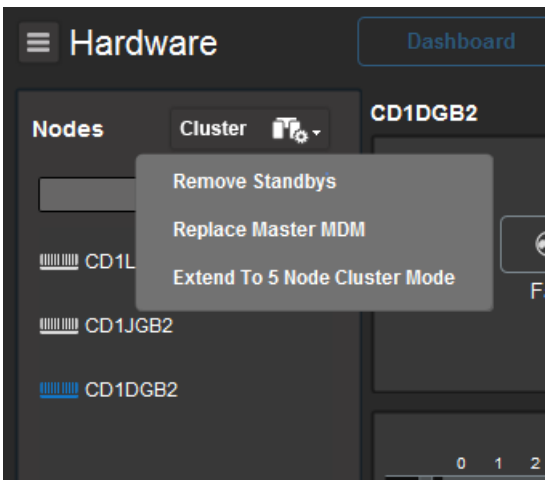

The **Switch to 5 Node Cluster Mode** window is displayed.

2. Select one Standby MDM and one Standby Tie Breaker.
3. Click **OK**.

Results

The selected standbys are added to the cluster, and the cluster is now operating in 5-node mode.

Switching from 5-node cluster mode to 3-node cluster mode

Decrease the number of nodes in an MDM cluster from five to three nodes, using the VxFlex OS GUI.

When possible, it is preferable to work in 5-node mode, especially in large systems. However, if it is necessary to reduce the number of nodes participating in the MDM cluster, follow these procedures to ensure a graceful change of cluster size.

Procedure

1. In the **Hardware** view, click the **Cluster** menu, and select **Decrease to 3 Node Cluster Mode**.

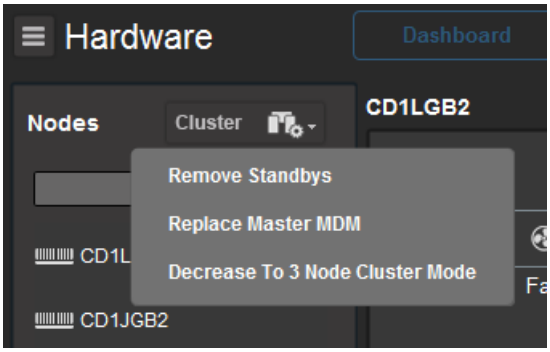

The **Switch to 3 Node Cluster Mode** window is displayed.

2. Select one MDM and one Tie Breaker for removal from the cluster.
3. Click **OK**.

Results

The cluster switches to 3-node cluster mode. The nodes that you removed from the MDM cluster revert to standbys.

Replacing a Master MDM

Replace a Master MDM with a Slave MDM, using the VxFlex OS GUI.

If you need to remove a server from a system, and that server is currently hosting the Master MDM, switch cluster control from the Master MDM to one of the Slave MDMs in the cluster. This operation ensures a graceful transfer of control, before removal of the server.

Procedure

1. In the **Hardware** view, click the **Cluster** menu, and select **Replace Master MDM**.

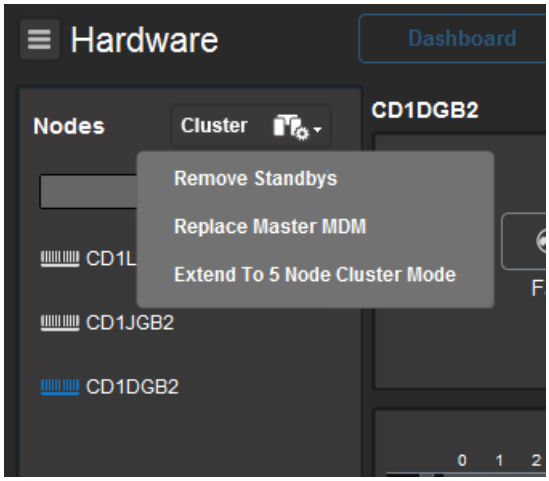

The **Replace Master MDM** window is displayed.

2. Select the Slave MDM to which you want to transfer system control.
3. Click **OK**.

Results

The selected Slave MDM becomes the Master MDM.

Transfer the MDM cluster component using the AMS

Switch an MDM or Tie Breaker (TB) component from one VxFlex Ready Node server to another.

Before you begin

Before moving the cluster component, both the source and the target servers must be part of the VxFlex OS system.

The source is the server from which the MDM or TB will be moved and the target is the server to which the MDM or TB will be moved. If there is no server in the system that can accept the role of an MDM or Tie Breaker component, you must add one before performing this procedure.

Procedure

1. If there is no standby node in the system, add a standby node:
   a. Use the VxFlex OS GUI Hardware view to find a server with a non-MDM role.

      This server will be the target. For example:

Figure 40 Target server

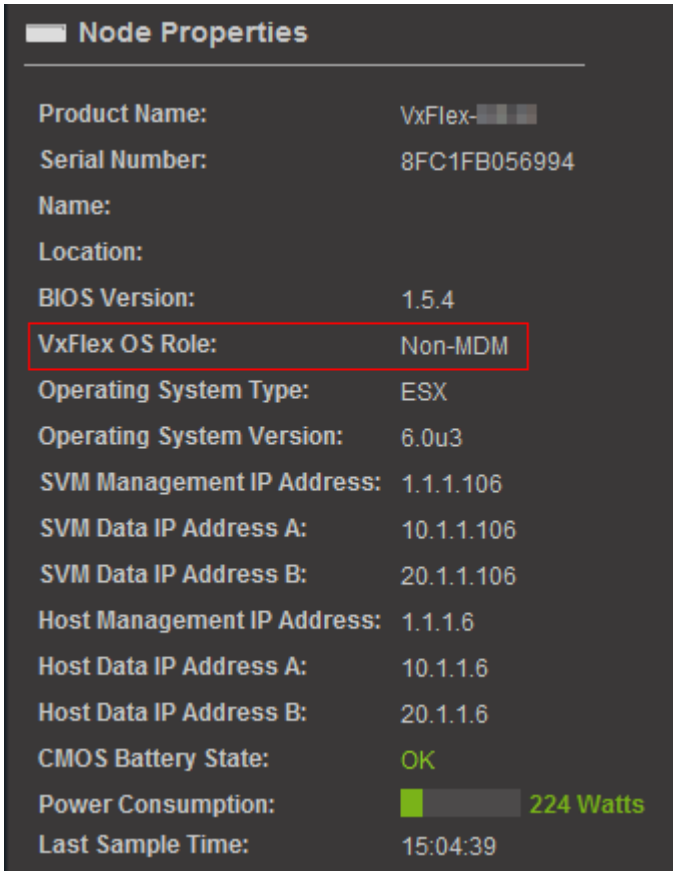

b. In **Hardware** view, right-click the target server and select **Add Standby Node**.

c. In the **Add Standby Node** dialog box, enter a name for the node, select the role that the source server had, and then click **OK**.

2. Right-click the target server and select **Replace Cluster Member**.
3. In the **Replace Cluster Member** dialog box, select the source server (the faulted server) and then click **OK**.
4. When the cluster component has been successfully transferred from the source to the target server, a success message is displayed. Click **Close**.

   The Node Properties sheet displays the source server as a standby MDM and the target server as an MDM in the cluster.

Figure 41 Source server, now a standby MDM

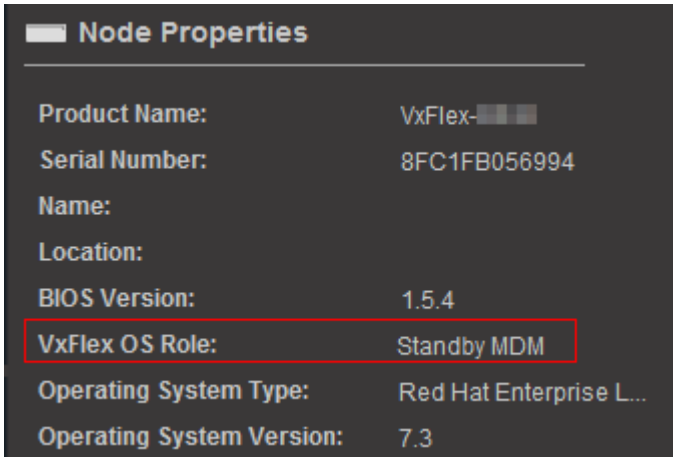

Figure 42 Target server, now an MDM

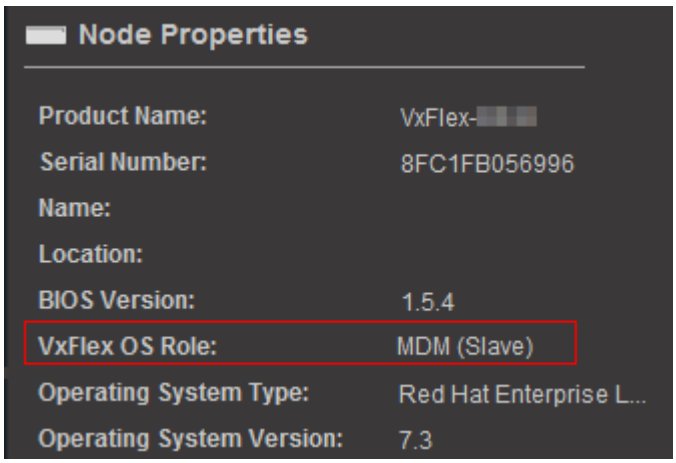

If the process was not successful, contact Customer Support. The virtual IP address on the MDM cluster ensures that all SDCs remain connected to the MDM, even after the physical server is replaced.

Configuring Oscillating Failure counters

Oscillating failure handling provides the ability to handle error situations, and to reduce their impact on normal system operation. This feature detects and reports various oscillating failures, in cases when components fail repeatedly and cause unnecessary failovers. You can configure the time interval associated with each window type, and the number of failures allowed before reporting commences for each window type, per counter.

You can reset specified oscillating failure counters to zero. This can be useful when you have fixed a problem and want to ensure that an alert is no longer active in the system.

Configuring Oscillating Failure counter parameters

Configure Oscillating Failure counter parameters for the entire system, or for specific Protection Domains or Storage Pools.

Procedure

1. Perform one of the following:

Option	Description
To configure counter parameters for all SDCs, Protection Domains or Storage Pools in the system:	In the **Backend** > **Storage** view, select the **System** icon.
To configure counter parameters for a specific Protection Domain or Storage Pool:	In the **Backend** > **Storage** view, navigate to, and select the desired Protection Domain or Storage Pool.

2. From the **Command menu** or context-sensitive menu, select **Set Oscillating Failure Properties**.
3. Perform one of the following:

Option	Description
For system level:	In the **For All** box, select an option: **SDCs**, **Protection Domains**, or**Storage Pools**.
For a Protection Domain or a Storage Pool:	Go to the next step.

4. In the **Counter Type** box, select a counter. Options vary, depending on the item selected in the previous step.
5. In the **Window Type** box, select an option for the sliding window interval: **Short**, **Medium** or **Long**.
6. Perform one of the following:

Option	Description
If you want to remove the selected counter definition from the system:	Select the **Remove the counter** check box.
If you want to modify the threshold for the selected counter definition:	Enter a number in the fields for: • **failures** (the maximum number of failures per time interval before reporting begins) • **seconds** (the number of seconds per time interval)

7. Click **OK**.

   The currently configured counter parameters are displayed in the corresponding Property Sheet, in the **Oscillating Failure Parameters** section.

Resetting Oscillating Failure counters

You can reset specified oscillating failure counters to zero. This can be useful when you have fixed a problem and want to ensure that an alert is no longer active in the

system. You can reset counters for the entire system, per Protection Domain, or per Storage Pool.

- To reset oscillating failure counters for all SDCs, Protection Domains or Storage Pools in the system, go to step 1.
- To reset counters for a specific Protection Domain or Storage Pool, go to step 6 on page 214.

Procedure

1. In the **Backend** > **Storage** view, select the **System** icon.
2. Right-click and select **Reset Oscillating Failure Counters**.
3. In the **For All** box, select an option.

   If you want to reset counters for all object types, select **Objects**.
4. Perform one of the following:
   - For a specific counter, in the **Counter Type** box, select the required counter.
   - For all counters, in the **Counter Type** box, select **None**, and select the **Reset All Counters** check box.
5. Click **OK**.

   To reset counters for a specific Protection Domain or Storage Pool, perform the following steps:
6. In the **Backend** > **Storage** view, navigate to, and select the desired Protection Domain or Storage Pool.
7. Right-click and select **Reset Oscillating Failure Counters**.
8. Perform one of the following:
   - For a specific counter, in the **Counter Type** box, select the required counter.
   - For all counters, select the **Reset All Counters** check box.
9. Click **OK**.

Viewing Oscillating Failure counters

View Oscillating Failure counters for network related issues, for SDCs, for SDSs and for devices.

You can view Oscillating Failure counters for network related issues, for SDCs, for SDSs and for devices in the following ways:

- Network:

1. In the **Backend** > **Storage** view, select the **System** icon.
2. From the **Command menu** or context-sensitive menu, select **Download Network Failure Counters**.
3. Browse to the location in which the file will be saved, and click **OK**. A JSON file containing the counters is saved in the location that you specified.
   - In Windows, view the file in a text editor, such as Notepad++.
   - In Linux, use the `more` command to view the file (for example, `more Oscillating_Network_Failures_Counters_2016-01-28-13-31-57.json`)

- SDCs:

1. In the **Frontend** > **SDCs** view, select the required SDC.
2. Open the Property Sheet, and click the **Oscillating Failure Counters** section.

Note

When there are no oscillating failures counters for SDC, the **Oscillating Failures Counters** section displays `None Found`.

- SDSs and devices:

Procedure

1. In the **Backend** > **Storage** view, navigate to and select the required SDS or device.
2. Open the Property Sheet, and click the **Oscillating Failure Counters** section.
   a. If there is an active alert, you can view the oscillating failures alert in the **Alerts** section of the Property Sheet.

Note

When there are no oscillating failures counters for SDS or device, the **Oscillating Failures Counters** section displays `None Found`.

Applying Performance Profiles to system components

You can use the VxFlex OS GUI to apply performance profiles to system components. The high performance profile configures a predefined set of parameters for very high performance use cases. When a container is provided in the command (System/Protection Domain/Fault Set), all the objects currently in the container are configured.

Note

For a complete list of parameters controlled by the profiles, refer to the document *VxFlex OS Performance Fine-Tuning Technical Notes*.

The profiles are applied separately to:

- SDSs
- SDCs
- MDM cluster

Note

After changing the performance profile of an SDS (on an SVM), you must perform manual memory allocation on the SVM, as described in the *VxFlex OS Deployment Guide*.

To apply a profile to system components, perform the following steps:

Procedure

1. Depending on the system component that you want to configure, in either the **Backend** > **Storage** or **Frontend** > **SDCs** view, navigate to, and select the desired objects.

Note

If you want to apply the Performance Profile to MDMs, select the System object.

2. Right-click the object and select **Set Performance Profile for XXX**, where XXX represents one of the following:
   - MDMs
   - All SDSs
   - SDS
   - All SDCs
   - SDC

   The **Set Performance Profile for** window is displayed.
3. Select a profile from the drop-down list, and click **OK**.

Configuring I/O priorities and bandwidth use (advanced)

VxFlex OS includes advanced settings which control I/O priorities and bandwidth use, which can be used to fine-tune system performance. It is recommended to retain default settings, unless you are an advanced user.

Application IOPS and bandwidth (advanced)

Priority can be given to different types of I/Os in the system, including application I/Os. The number of concurrent Rebuild and Rebalance jobs can be configured, along with bandwidth used for these jobs. I/O prioritization is configured per Storage Pool.

NOTICE

These features affect system performance, and should only be configured by an advanced user.

Give priority to Application I/Os during Rebuild and Rebalance jobs, by performing these steps:

Procedure

1. In the **Backend** > **Storage**view, navigate to, and select the desired Storage Pool.
2. Right-click the Storage Pool and select **Set I/O Priority**.
3. Select **Favor Application I/O for Rebalance and Rebuild**, and click **OK**.

System IOPS and bandwidth (advanced)

You can prioritize different types of I/O in the system, as well as configure network throttling per Protection Domain.

I/O prioritization

Priority can be given to different types of I/Os in the system. The number of concurrent Rebuild and Rebalance jobs can be configured, and bandwidth for Rebalance jobs can be configured. If the `Dynamic Bandwidth Throttling` option

is selected, additional items can be configured, such as `Application IOPS threshold`, `Application bandwidth threshold`, and `Application threshold quiet period`. Default values for these features are provided in the *VxFlex OS CLI Reference Guide*.

NOTICE

These features affect system performance, and should only be configured by an advanced user.

Configure I/O prioritization for Rebuild and Rebalance by performing these steps:

Procedure

1. In the **Backend** > **Storage** view, navigate to, and select the desired Storage Pool.
2. Right-click the Storage Pool and select **Set I/O Priority**.
3. Select the desired options and edit values, and click **OK**.

Configuring Network Throttling

Configure network throttling to control the flow of traffic over the network.

Network throttling is configured per Protection Domain. The SDS nodes transfer data between themselves. This data consists of user-data being replicated as part of the RAID protection, and data copied for internal rebalancing and recovery from failures. You can modify the balance between these types of data loads by limiting the data copy bandwidth. This change affects all SDSs in the specified Protection Domain.

NOTICE

These features affect system performance, and should only be configured by an advanced user. Contact customer support before you change this configuration.

Procedure

1. In the **Backend** > **Storage** view, navigate to the desired Protection Domain, and select its row in the table.
2. Right-click the row and select **Set Network Throttling**.

   The **Set Network Throttling** window is displayed.
3. Configure the settings, and click **OK**.

Enable and disable Rebuild/Rebalance (advanced)

By default, Rebuild and Rebalance features are enabled in the system, because they are essential for system health, optimal performance, and data protection. These features should only be disabled temporarily in very specific circumstances, and should not be left disabled for long periods of time. Rebuild and Rebalance features are enabled and disabled per Storage Pool.

NOTICE

Rebuilding is an essential part of VxFlex OS, which provides protection for your data. It is not recommended to disable the Rebuild feature, except in very special circumstances. Rebalancing is an essential part of VxFlex OS, and should only be disabled, temporarily, in special circumstances. Disabling rebalance may cause the system to become unbalanced even if no capacity is added or removed. For example, during a recovery from an SDS or device failure, some rebalance activity may be needed to ensure optimal balancing.

To enable or disable Rebuild and Rebalance features, perform these steps:

Procedure

1. In the **Backend** > **Storage** view, navigate to, and select the desired Storage Pools.
2. Right-click the Storage Pool and select **Enable/Disable Rebuild/Rebalance**.

   The **Enable or Disable Rebuild and Rebalance** window is displayed.
3. Select or clear the options that you require (selected=enable; clear=disable), and click **OK**.

Using the background device scanner

The background device scanner scans devices in the system to check for errors. You can enable and disable the background device scanner, as well as reset the background device scanner counters. Information about errors will be provided in event reports. For more information about viewing events, see Viewing events on page 265.

Enabling and disabling the background device scanner

The scanner can be enabled on all the devices in the specified Storage Pool. There are two modes: **device only** mode, and **data comparison** mode:

- Device only—Perform read operations. Fix from peer on errors.
- Data comparison—Perform the device-only test, and compare the data content with peer. Zero padding must be enabled in order to set the background device scanner to data comparison mode.

  To enable or disable the background device scanner, follow these steps:

Procedure

1. In the **Backend** > **Storage** view, navigate to, and select the desired Storage Pools.
2. From the **Command menu** or context-sensitive menu, select **Set Background Device Scanner Mode**.

   The **Configure Storage Pool Background Device Scanner** window is displayed. The right pane of the window displays the Storage Pools that you are configuring.
3. For the **Enable Background Device Scanner** option, perform one of the following:
   - To enable the scanner, select the check box, and proceed to the next step.

- To disable the scanner, clear the check box, and click **OK** to finish.

4. Select an option (selected=enable; clear=disable):
   a. **Device only**
   b. **Data comparison**
5. In the **Bandwidth Limit** box, accept the default or type a number in KB per second (per device).

   The given value should be in the range 10 KB-10 MB (default = 1 MB).

   Note

   High bandwidth should be used very carefully for extreme cases only (such as an urgent need to check some devices), as it may create negative impact on system performance. Setting the background device scanner bandwidth should take into account maximum bandwidth of the devices.

6. Click **OK**.

Resetting the background device scanner counters

You can reset background device scanner error counters for specified Storage Pools. Counters for data comparison errors, or corrected read errors, or both counter types can be reset.

To reset counters, follow these steps:

Procedure

1. In the **Backend** > **Storage** view, navigate to, and select the desired Storage Pools.
2. Right-click the Storage Pools and select **Reset Background Device Scanner Counters**.

   The **Reset Background Device Scanner Counters** window is displayed. The right pane of the window shows the Storage Pools that you are configuring.
3. Select or clear the option that you require, or both options (selected=enable; clear=disable).
4. Click **OK**.

Modifying Checksum protection mode

Checksum mode can be used to validate in-flight data reads and writes, in order to protect data from data corruption. To modify this setting, perform the following steps:

Procedure

1. In the **Backend** > **Storage** view, navigate to, and select the desired Storage Pools.
2. From the **Command menu** or context-sensitive menu, select **Configure Use Checksum**.

   The **Configure Use Checksum** window is displayed.
3. Do one of the following:
   a. To enable the Checksum feature, select the **Enable Use Checksum** check box.

b. To disable the Checksum feature, clear the **Enable Use Checksum** check box.

4. Click **OK**.

Renaming objects

Object names are used to identify the objects in the VxFlex OS GUI, and can also be used to specify objects in CLI commands. You can view an object's name in its Property Sheet, in the **Identity** section.

Note

It is not possible to rename a Read Flash Cache device using this command.

You can define object names according to the following rules:

1. Contain less than 32 characters
2. Contain only alphanumeric and punctuation characters
3. Be unique within the object type

When a name has not been defined, the system may display default system-defined names, as follows:

- SDC—its first IP address
- SDS—its first IP address
- Device—the path to the device
- All other objects—the object's ID

Note

A name must be assigned to a volume when it is initially created. You can rename the volume later, using the **Rename** command.

Procedure

1. Depending on the object type, in the **Backend** > **Storage** or any of the **Frontend** views, navigate to the object in the table, and select its row.
2. Right-click the object and select **Rename**.

   An editing window is displayed, showing the current name, and an editable field for the new name.
3. Type the new name in the field, and click **OK**.

Provision compute, storage and network environments, and deploy workloads

Compute, storage and network environments can be provisioned and managed from within the VxFlex Ready Node system. Workloads can then be deployed in the environments. Predefined profiles simplify these procedures. This entire feature is optional.

Prerequisites:

- vCenter credentials must be configured in the VxFlex Ready Node system (**System Settings** > **vCenter**).

- ESX servers must already be managed by (visible in) the vCenter.
- The vCenter must contain one or more VM templates, which are configured for your requirements.
- IP address pools must be configured in your VxFlex Ready Node system (**System Settings** > **VMkernel IP Pools**).
- If you plan to use a Virtual Distributed Switch in Active-Active or Active-Standby mode, configure MLAG (Multichassis Link Aggregation) in the physical switch. Otherwise, the network connection will not work in the vSphere Distributed Switch.
- If you plan to use a vSwitch in Active-Standby mode, configure MLAG (Multichassis Link Aggregation) in the physical switch. Otherwise, the network connection will not work in the vSphere Distributed Switch.
- If a separate migration network is desired, define a migration network first, so that it is available for selection when performing this procedure. In vSphere, create a vSwitch, and add the nodes. Do not add nodes to sio_vswitch_1 and sio_vswitch_2.
- In order to be able to use the VxFlex OS data network as the migration network in the compute environment, you should enable VMotion on VxFlex OS data NICs in the vCenter.
- Ensure that the MTU value for VMkernel and network interfaces are configured identically in the VxFlex OS GUI and in vCenter.

The information displayed in the **Compute** view reflect the clusters and workloads displayed in the vCenter.

The basic workflow is:

1. Create VMkernel IP address pools (**System Settings** > **VMkernel IP Pools**).
2. Create a parent environment: Use default settings, existing profiles or manually configure network and compute environments. Configure storage settings.
3. Create a workload and deploy it.
4. Use the toolbar buttons in the **Compute** view to perform various monitoring or modification tasks, or to create a new environment. For more information about monitoring compute clusters and workloads, see Monitoring compute and network environments on page 176.

Note

Several steps in the procedures can by performed using default profiles and settings. Compute, storage and network environment default values on page 232 lists the default values.

Configure VMkernel IP address pools

VMkernel IP address pools are used for configuring the VMkernel on the vCenter server that will be used for vMotion (moving VMs between hosts). The IP address pools are also used to check the connectivity of the network that is generated in creation of a compute environment. Each IP address pool contains a consecutive range of valid IP addresses. Create one large range (pool) or several ranges (pools).

Before you begin

Ensure that you have enough IP addresses for all your hosts.

You can add, modify and delete VMkernel IP address pools.

Procedure

1. From the **System Settings** menu at the top right side of the window, in any view, select the **System Settings** option.

   The **System Settings** window is displayed.

2. In the left panel, click **VMkernel IP Pools**.

   The **Application Network VMkernel IP Pool** panel is displayed on the right side of the window.

3. To configure IP address pools, do one of the following:

Option	Procedure
Create an IP address pool	Click the **Add IP Pool** button. In the lower part of the window, type a name for the IP pool in the **Name** box, a subnet mask in the **Subnet Mask** box, and IP addresses in the **IP Address Ranges** boxes. Click the **+** button to add more IP ranges. Click **Save**. **Note** You can clear the **Add IP Pool** section and create another IP Pool by clicking the **X** in the **Add IP Pool** section.
Modify an IP address pool	Right-click the pool that you want to modify, and select **Update IP Pool**, or select the pool and click the **Update IP Pool** button. Modify the items displayed in the lower part of the window, and click **Save** when you have finished.
Remove an IP address pool	Right-click the pool that you want to delete, and select **Remove IP Pool**. Click **Save**.

Provision an environment in simple mode

Take advantage of built-in defaults and system recommendations to provision a compute, storage and network environment.

Before you begin

- vCenter credentials must be configured in the VxFlex Ready Node system (**System Settings** > **vCenter**).
- ESX servers must already be managed by (visible in) the vCenter.
- The vCenter must contain one or more VM templates, which are configured for your requirements.
- IP address pools must be configured in your VxFlex Ready Node system (**System Settings** > **VMkernel IP Pools**).
- If you plan to use a Virtual Distributed Switch in Active-Active or Active-Standby mode, configure MLAG (Multichassis Link Aggregation) in the physical switch. Otherwise, the network connection will not work in the vSphere Distributed Switch.

- If you plan to use a vSwitch in Active-Standby mode, configure MLAG (Multichassis Link Aggregation) in the physical switch. Otherwise, the network connection will not work in the vSphere Distributed Switch.
- If a separate migration network is desired, define a migration network first, so that it is available for selection when performing this procedure. In vSphere, create a vSwitch, and add the nodes. Do not add nodes to sio_vswitch_1 and sio_vswitch_2.
- In order to be able to use the VxFlex OS data network as the migration network in the compute environment, you should enable VMotion on VxFlex OS data NICs in the vCenter.
- Ensure that the MTU value for VMkernel and network interfaces are configured identically in the VxFlex OS GUI and in vCenter.

You can provision an environment in simple mode or in advanced mode. Simple mode uses auto-select functionality, which suggests recommended settings. An environment profile may be used (default, or customized previously). This profile will automatically populate most settings. Compute, storage and network environment default values on page 232 lists the default values.

Procedure

1. Click the **Frontend** tab, and select **Compute**.

   The **Compute** view appears.
2. If an explanatory window appears, read the information and then close it.

   You can refer back to this information at any time, by clicking the **Help** button in the **Compute** view's toolbar.
3. Click the **+** button.
4. Select a profile.
5. In the **Environment Name** box, type a name.
6. In the **Compute** section, type the desired number of nodes (two or more nodes), and then perform one of the following:
   - To allow the system to select the nodes automatically, click the **Auto-Select** button next to the number of nodes.
   - To select the nodes manually, click the **Edit selection** button near the number of nodes, and select the desired servers from the list in the right pane of the window.
7. By default, the Storage Pool with the most free capacity is selected. To choose a different Storage Pool, click **Select Storage Pool**, and select one or more Storage Pools listed in the right pane of the window.
8. In the **Network** section, **# of Switches** box, select the number of switches to be used in the environment.

   The maximum number allowed is 3.
9. In the **Network** section, expand the **Advanced** section, accept the default profile, or select a different profile.
10. When all configurations are complete, click **Create Environment**.

    In the right pane of the window, progress is displayed. The **Command Monitor** button on the toolbar displays the commands currently being executed, and the result of each command.
    Once the environment is created, the table in the main part of the **Compute** view shows the system's environment resources. The **Environment** blade

remains in read only mode, displaying the spec of the environment that was created. The following example shows clusters in the VxFlex OS GUI and in vCenter.

Figure 43 Clusters shown in VxFlex OS GUI and in vCenter

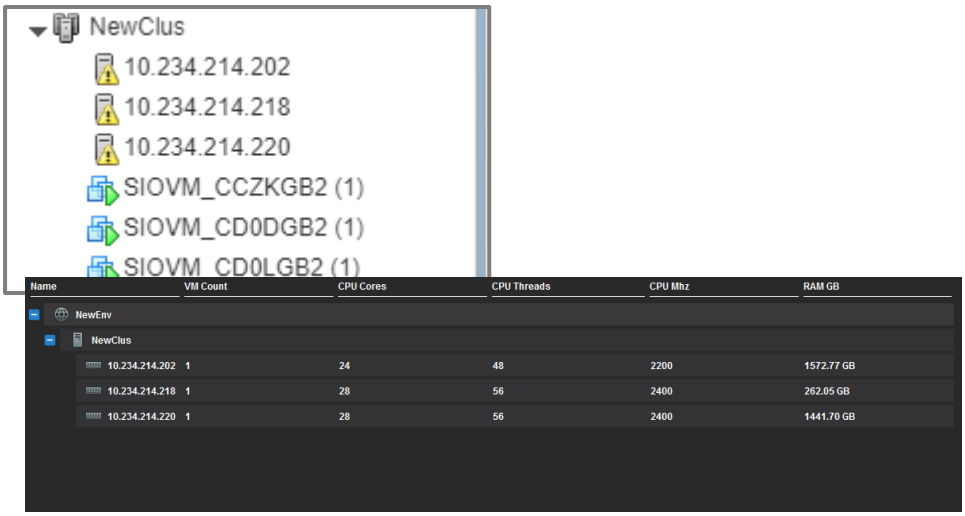

You can now proceed to deploy a workload in the environment.

Provision an environment in advanced mode

Use profiles or manually provision a compute, storage and network environment. Predefined profiles simplify these procedures.

Before you begin

- vCenter credentials must be configured in the VxFlex Ready Node system (**System Settings** > **vCenter**).
- ESX servers must already be managed by (visible in) the vCenter.
- The vCenter must contain one or more VM templates, which are configured for your requirements.
- IP address pools must be configured in your VxFlex Ready Node system (**System Settings** > **VMkernel IP Pools**).
- If you plan to use a Virtual Distributed Switch in Active-Active or Active-Standby mode, configure MLAG (Multichassis Link Aggregation) in the physical switch. Otherwise, the network connection will not work in the vSphere Distributed Switch.
- If you plan to use a vSwitch in Active-Standby mode, configure MLAG (Multichassis Link Aggregation) in the physical switch. Otherwise, the network connection will not work in the vSphere Distributed Switch.
- If a separate migration network is desired, define a migration network first, so that it is available for selection when performing this procedure. In vSphere, create a vSwitch, and add the nodes. Do not add nodes to sio_vswitch_1 and sio_vswitch_2.
- In order to be able to use the VxFlex OS data network as the migration network in the compute environment, you should enable VMotion on VxFlex OS data NICs in the vCenter.
- Ensure that the MTU value for VMkernel and network interfaces are configured identically in the VxFlex OS GUI and in vCenter.

You can provision an environment in simple mode or in advanced mode. In advanced mode, several steps in the procedure can by performed using profiles (default or

customized) and settings. Compute, storage and network environment default values on page 232 lists the default values.

Note

You can create customized profiles for future use by clicking the **Save** buttons next to each of the sections that you configure. Alternatively, configure an entire environment and then click the **Save** button next to the Environment name. An environment profile will be saved, and the network and compute profiles will be created automatically.

Note

Do not change the names of clusters or NICs on VMware after using the VxFlex Ready Node to set up your compute environment.

Procedure

1. Click the **Frontend** tab, and select **Compute**.

   The **Compute** view appears.
2. If an explanatory window appears, read the information and then close it.

   You can refer back to this information at any time, by clicking the **Help** button in the **Compute** view's toolbar.
3. On the toolbar in this view, click the **+** button.

   The **Environment** pane appears.
4. Ensure that the **Profile** box is empty.
5. In the **Environment Name** box, type a name.
6. In the **Compute** section, type the desired number of nodes, and then perform one of the following:
   - To allow the system to select the nodes automatically, click the **Auto-Select** button next to the number of nodes.
   - To select the nodes manually, click the **Edit selection** button near the number of nodes, and select the desired servers from the list in the right pane of the window.
7. Expand the **Advanced** section, and either select a compute profile, or clear the **Profile** box and configure all the fields in this section.
   a. In the **# of Clusters** box, type the number of clusters that you wish to create.
   b. In the **Cluster Prefix** box, type text.

      This text will be added to the name of each cluster in the series that is created.
   c. In the **HA** box, select the checkbox to enable High Availability, or clear it to disable High Availability.
   d. In the **DRS** box, select the checkbox to enable DRS (VMware Distributed Resource Scheduler), or clear it to disable DRS.
   e. If you wish to use the Overcommit feature, select the **Overcommit** checkbox, and type the percentage of overcommit to apply to the clusters, for **Total Logical CPUs** and for **Total RAM**.

This feature configures the system for more resources than actually exist for CPU and RAM, as a percentage. The maximum percentage allowed is 500% for RAM and 200% for CPU.

8. By default, the Storage Pool with the most free capacity is selected. To choose a different Storage Pool, click **Select Storage Pool**, and select one or more Storage Pools listed in the right pane of the window.
9. Expand the **Advanced** section, and configure all the fields in this section.
   a. Type the size of storage required for each cluster, and select a unit type.
   b. In the **Datastore Prefix** box, type text.

      The text will be added to the name of each datastore in the series that is created.

   c. If you wish to use the Overcommit feature, select the **Overcommit** checkbox, and type the percentage of overcommit to apply to the storage.
10. In the **Network** section, **# of Switches** box, select the number of switches to be used in the environment.
11. Expand the **Advanced** section and perform one of the following:
    - If you intend to create a profile or perform manual configurations, clear the **Profile** box
    - Select the default profile
    - Select a custom profile
12. Configure Virtual Switches:
    a. In the **Name** column, select a name, or type a new name.
    b. In the **Migration** column, select or clear the checkbox depending on whether or not the Virtual Switch will be used for VM migration.

       When checked, the virtual switch will support VM migration. When cleared, you must select a network for migration in the **Migration network** box.

    c. Select the number of VM networks to be supported by the Virtual Switch. Up to 10 networks can be supported.
    d. In the **Prefix** box, type text representing the prefix that will be appended to the VM network name.

       The text will be added to the name of each VM network in the series that is created.

    e. If you did not select the **Migration** option, select a different network for VM migration in the **Migration network** box.
    f. In the **NIC** column, click **+** to select the NIC of each server to be assigned to the VM network.

       Optional: Right-click on nodes and apply the same NIC selection on the matching nodes.

       The default settings are: VxFlex OS Data and Management networks.

Note

If one of the VxFlex OS NICs is selected, you can not select other NICs for VMotion. If client NICs are selected for VMotion, all VxFlex OS options will disappear.

Note

If one of the NICs is already configured in the vCenter, you can use it for the environment that you are configuring now. It will be presented as Homogeneously Configured Network in the VxFlex OS GUI. However, if this option is used, you cannot modify any of the settings using the VxFlex OS GUI.

Note

The default profile enables Active-Active NIC link aggregation, and enables both Client data ports. However, this option is only supported if a VDS license exists in the vCenter.

13. Expand the **VM Network and Switch Settings** section, click the tabs for each VM network, and configure the following for each network:

    Note

    The VxFlex OS data and management networks cannot be used for the VM network.

    a. For each VM network, configure **MTU** (range allowed is 1500-9000) and **VLAN** (when used).

    b. Select a **Switch Type** option.

    c. Select a **Link Aggregation** mode.

    d. In the **VMKernel** section, select one or more IP pools, and configure their **MTU** and **VLAN** values.

14. At this point, you can create profiles for the configurations that you have performed. Navigate to each section for which you would like to create a profile, and click the corresponding **Save** button.

15. When all configurations are complete, click **Create Environment**.

    In the right pane of the window, progress is displayed. The **Command Monitor** button on the toolbar displays the commands currently being executed, and the result of each command.
    Once the environment is created, the table in the main part of the **Compute** view shows the system's environment resources. The **Environment** blade remains in read only mode, displaying the spec of the environment that was created. The following example shows clusters in the VxFlex OS GUI and in vCenter.

Figure 44 Clusters shown in VxFlex OS GUI and in vCenter

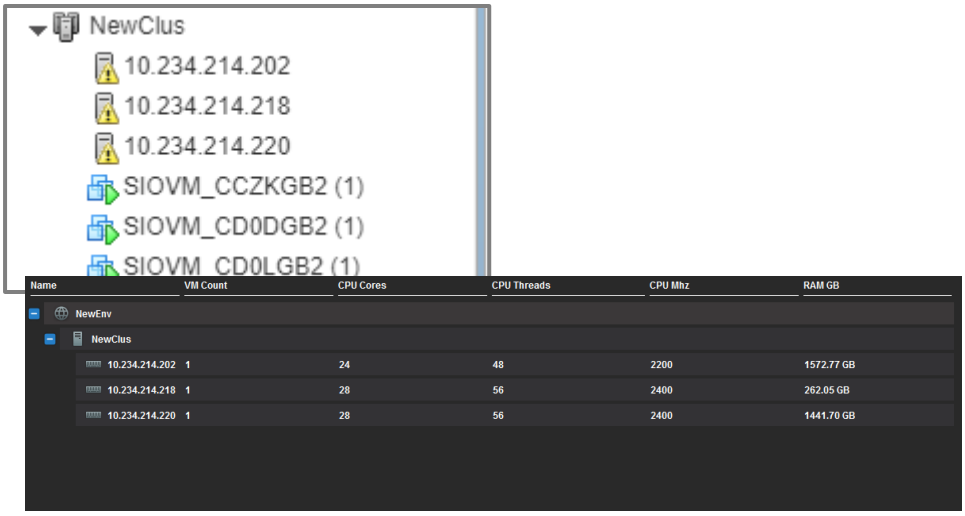

You can now proceed to deploy a workload in the environment.

Deploy a workload

Deploy one or more workloads in compute environments that have already been provisioned using the **Compute** view.

Before you begin

- vCenter credentials must be configured in the VxFlex Ready Node system (**System Settings** > **vCenter**).
- ESX servers must already be managed by (visible in) the vCenter.
- The vCenter must contain one or more VM templates, which are configured for your requirements.
- IP address pools must be configured in your VxFlex Ready Node system (**System Settings** > **VMkernel IP Pools**).
- At least one Environment should exist in the system.

Procedure

1. Click the **Frontend** tab, and select **Compute**.

   The **Compute** view appears.
2. Click the **Environment Resources** toolbar button.
3. In the main table, right-click an environment, and select **Deploy Workload**.

   The **Workload** pane appears.
4. In the main pane of the window, select a VM template.
5. For Windows deployments, in the **Workload** pane,**Windows License** box, type the Windows license.
6. In the **Workload** pane, in the **VM Name** box, type a name.
7. In the **No. of VMs** box, type the number of VMs required for the workload.
8. To accept the default settings from the VM template, proceed to the next step. For advanced configurations, expand the **Advanced** section and modify the following settings as required.

- Virtual sockets
- Cores per socket
- RAM

Note

Changes made to the advance settings have no effect on the template in the vCenter.
If you do not want to configure workload IP addresses, or Windows license, or if the template does not have VMware Tools, do not select the **Customize VMs** checkbox.

9. In the **Adapters** section, click a NIC, and then the **Edit** button.

   The **Workload Network Settings** window appears for the NIC.
10. Select a protocol.
    - If you choose **Static**, configure one or more IP address ranges, the subnet mask and the gateway to be used by the VMs.
    - If you choose **DHCP**, configure a subnet mask and a gateway address.
11. Select a network.

    The Network list displays all the networks configured in the vCenter.
12. Click **Save**

    The configurations will be saved on the VMs that are added to the workload. In the left pane, you can now deploy the workload.
13. Click **Deploy**.

Results

In the right pane of the window, progress is displayed. The **Command Monitor** button on the toolbar displays the commands currently being executed, and the result of each command. Once the workload is deployed, it is visible in both the VxFlex OS GUI and in the vCenter, as shown in the example below.
Later, you can expand or delete workloads, and expand or delete environments, by right-clicking the entities in the main table in the **Compute** view.

Note

After the VMs are deployed, a "MAC Conflict" alert might be displayed in the vCenter on some of the VMs. This is a false alert, and you can either ignore or acknowledge it. For more information about this alert, see the following VMware Knowledge Base article: https://kb.vmware.com/s/article/1024025

Note

If issues are encountered when deploying Windows VMs, see the following VMware Knowledge Base article: https://kb.vmware.com/s/article/1025707

Figure 45 Workloads shown in VxFlex OS GUI and in vCenter

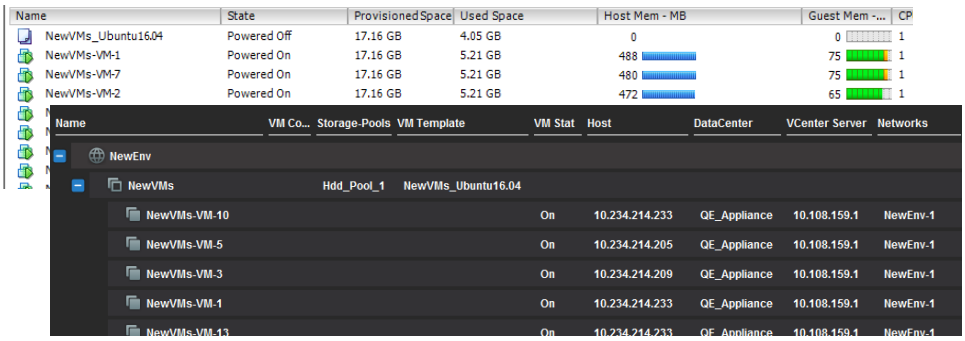

Expand environments

You can expand an existing environment by adding hosts to an existing cluster, without having to configure everything again from scratch. If you want to add a new cluster, when you select the **New** option, all the settings that were configured for the for the first compute environment are applied to the new cluster. These setting cannot be edited.

Note

If you have upgraded your system from version 2.0.1.4, the network profiles used may need to be edited. If no names are configured for vSwitches and VM Networks, add them during this expansion procedure.

Procedure

1. Click the **Frontend** tab, and select **Compute**.

   The **Compute** view appears.
2. In the main table, right-click an environment, and select **Expand Environment**.

   The **Select Nodes** and the **Environment** panes appear.
3. Select nodes to add to the cluster.
4. Scroll down to the **Network** section, and for the required virtual switch, click the **+** button in the **NIC** column.

   The Select NICs pane appears.
5. Expand the nodes to reveal the list of available NICs. Select the same number of NICs with the same speed as the other nodes in the environment.
6. In the **Compute** section, expand the **Advanced** section. For **Cluster Mode**, select the **New** or **Existing** option, and configure the remaining settings that are available for configuration.
   - For an existing cluster, click **Expand Environment**.
   - For a new cluster, type a name for the new cluster and then click **Expand Environment**.

Results

The hosts are added to the environment.

Delete environments

Delete existing environments using the VxFlex OS GUI **Compute** view.

Procedure

1. Click the **Frontend** tab, and select **Compute**.

   The **Compute** view appears.
2. Right-click the environment that you want to delete, and select **Delete Environment**.

   The **Delete Environment** window appears.
3. Click **Delete Environment**.

Remove a node from a compute cluster

Remove a node from a compute cluster.

Procedure

1. Click the **Frontend** tab, and select **Compute**.

   The **Compute** view appears.
2. On the toolbar, click **Environment Resources** .
3. In the table, expand an environment, right-click a node and select **Remove Node**.

   The **Remove Node** window appears.
4. Click **Remove**.

   The node is removed from the compute cluster.

Expand a workload

Add VMs to an existing workload using the VxFlex OS GUI **Compute** view.

The workload's settings are automatically applied to the VMs being added to the system. The only settings required are the number of VMs to be added, and NIC options. If static IP addresses are required, these need to be configured as well.

Procedure

1. Click the **Frontend** tab, and select **Compute**.

   The **Compute** view appears.
2. On the toolbar, click the **VMs per Workload** button.
3. Select and right-click one or more workload rows, and select **Expand Workload**.

   The **Expand Workload** window appears.
4. In the **Number of VMs** box, type the number of VMs that you want to add to the workload.
5. In the **Adapters** section, select either the **DHCP** or **Static** option.
6. If you selected the **Static** option, configure the IP address, Subnet Mask and Gateway details.

7. Click **Expand**.

Results

The VMs are added to the workloads. You can follow progress by clicking the **Command Monitor** button.

Delete a workload

Delete a workload using the **Compute** view in the VxFlex OS GUI.

Procedure

1. Click the **Frontend** tab, and select **Compute**.

   The **Compute** view appears.
2. On the toolbar, click the **VMs per Workload** button.
3. Right-click the workload that you want to delete, and select **Delete Workload**.

   The **Delete Workload** window appears.
4. Optionally, to power off the VMs, select the **Power off the VMs** option.

   If this option is not selected, the all the VMs must be shutdown manually. Otherwise, the deletion process will fail.
5. Click **Delete Workload**.

   The **Command Monitor** view displays progress of the deletion process. Click the **Command Monitor** toolbar button to view deletion status.

Compute, storage and network environment default values

To ease deployment and provisioning activities, several settings can be configured using default profiles and settings. The following tables list the default settings.

Compute

Item	Default Value	Range	Comment
Number of nodes	N/A	N/A	
Compute profile	Default Compute profile	N/A	
Number of clusters	1		
Cluster prefix	VxFlex OS Cluster	N/A	Text input
HA checkbox	Selected	Selected / Clear	
DRS checkbox	Selected	Selected / Clear	In the vCenter, the DRS (Distributed Resource Scheduler) configuration should be Fully Automated.
Sockets	N/A	N/A	Total # of sockets in all nodes. For example: 2 sockets * 3 nodes = 6
Total Logical CPUs	N/A	N/A	The total number of Logical cores per CPU

Item	Default Value	Range	Comment
			including the Hyper threading ability which doubles the physical cores
Overcommit	Selected	Selected / Clear Maximum overcommit for RAM is 200% , and 500% for CPU	This setting lets you create VMs that exceed the physical and logical resources of the node (CPU and RAM). In most cases, VMs do not use the full CPU and RAM resources, so a user familiar with the application and VM usage sizing can decide to allow deployment of more VMs than are possible without overcommit.

Storage

Item	Default	Range	Comment
Storage Pool	The Storage Pool with the most free space	N/A	
Size	None	N/A	Optional parameter that limits the minimum size of the storage allocated for environment, and ensures that Storage Pools that are too small are not selected
Datastore prefix	N/A	N/A	Free text input
Overcommit	100%	100-500%	Allows allocation of more storage for the datastore created than is actually needed, thus leaving space for VMs to grow

Network

Item	Default	Range	Comment
MTU	1500	1500-9000	
VLAN	No VLAN	1-4094	
Virtual Switch type	DVS	vSwitch / DVS	If the vCenter has a DVS license, the default is DVS. If not, the default is vSwitch.
Teaming mode	For one port: Single	Single / Active-	

Item	Default	Range	Comment
	For two ports and a DVS license: Active-Active For two ports without a DVS license, Active-Passive	Active / Active-Passive	
IP Pool	First free IP Pool from the list defined in System Settings		If the vMotion setting is not selected, the same IP pool is used for VMkernel and for IP pool, for testing connectivity
VMkernel MTU	1500	1500-9000	Only available if vMotion is selected
VMkernel VLAN	No VLAN	1-4094	Will cause the networks to grow according to number set here

Workload

Item	Default	Range	Comment
Snapshots used for workload creation	19	N/A	Workload creation is based on the VxFlex Ready Node's volume snapshot in the background. You can modify this by changing the value set for `maxSnapshotsInWorkload` in the `phoenix.properties` file located in: • Windows: `C:\Program Files\EMC\scaleio\AMS\webapps\ROOT\WEB-INF\classes` • Linux: `/opt/emc/scaleio/ams/webapps/ROOT/WEB-INF/classes/`

Approving pending security certificates

Approve pending security certificates, and view approved certificates in the **System Settings** window.

Note

When there are pending certificates in the system, they are listed in the **Backend > Storage** view > **State Summary** table, and in the **Monitor > Alerts** view.

Procedure

1. From the **System Settings** menu in the top right corner, select **System Settings**.

   The **System Settings** window appears, showing approved and pending certificates.

2. Scroll to connections that are **Pending Approval**, and expand the rows.
3. For each one, scroll to the bottom of the information about the required certificate, and click **Confirm**.

Customizing system preferences

You can customize various features in the VxFlex OS GUI using the **User Preferences** window. The following features can be customized:

- Refresh data rate
- Clear host history from previous sessions
- Calculation of I/O workload average rate shown on the Dashboard
- System clock display
- Advanced display mode for Dashboard, Backend internal views, Frontend internal views, and Property Sheet
- Log level

Procedure

1. From any location in the VxFlex OS GUI, open the **System Settings** menu in the top right corner, and select **User Preferences**.

   The **User Preferences** window is displayed.

2. Edit the options according to your needs, and click **Apply**.

Figure 46 User Preferences window

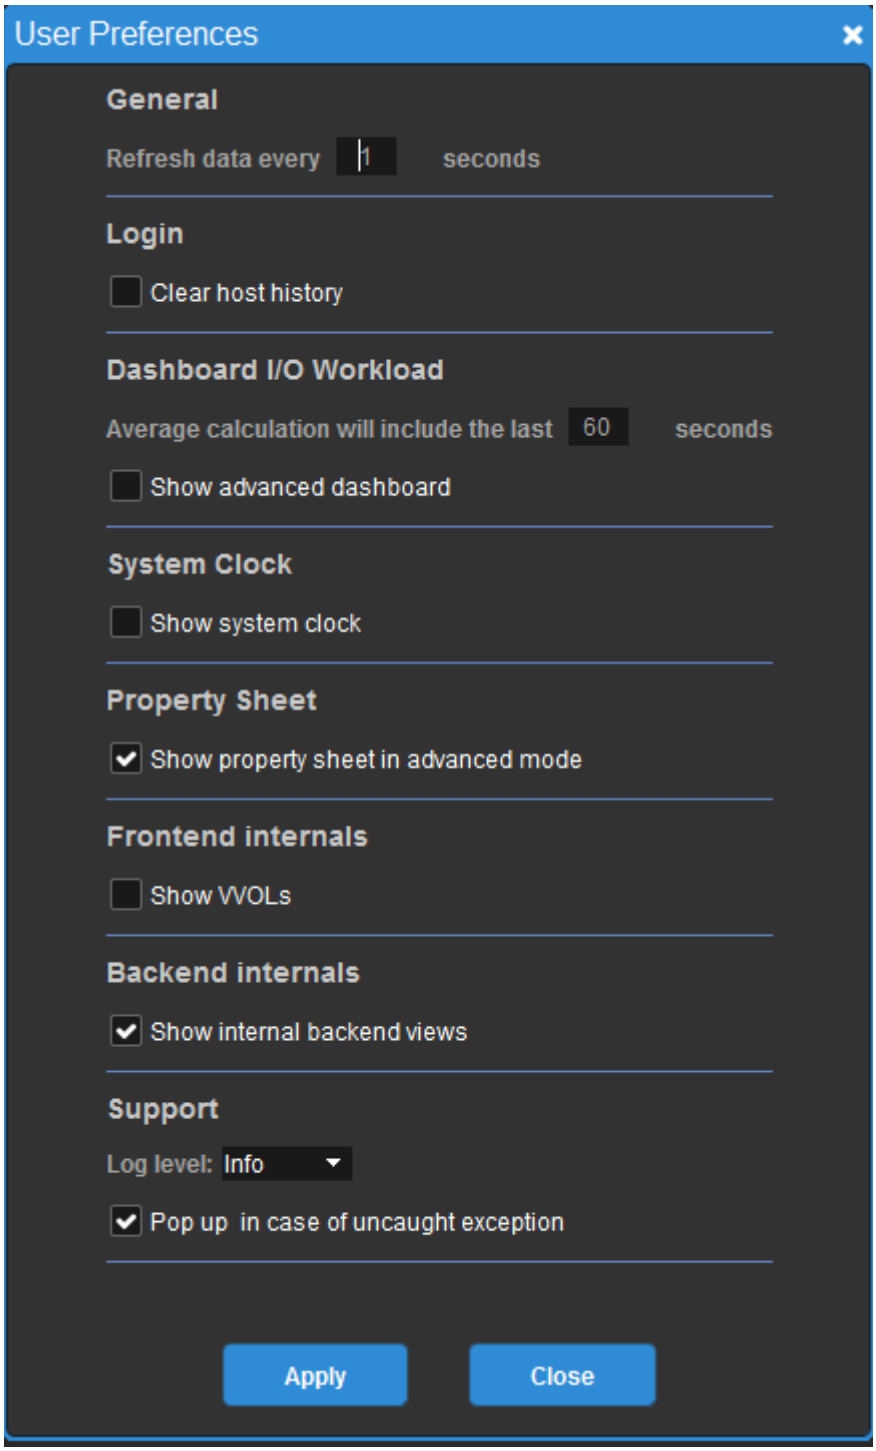

Table 31 User Preferences

Item	Description
General: Refresh data every *n* seconds	Controls the rate at which data displayed in the VxFlex OS GUI is refreshed, in seconds (Default: 10 seconds)

Table 31 User Preferences (continued)

Item	Description
	The refresh occurs at least at the specified rate. It is not intended to be used as a means of limiting client traffic, although it would actually do so.
Login: Clear host history	When selected, the VxFlex OS GUI does not save and present host connection details from previous sessions
Dashboard I/O workload: Average calculation will include the last *n* seconds	Controls the time period used when averages are computed and displayed by the VxFlex OS GUI (default: 10 seconds)
Show advanced dashboard	When selected, (default), includes more details on some tiles in the Dashboard view. The toggle buttons switch between the statistics displayed in large fonts and small fonts. The upper button toggles between average values and sample values. The lower button toggles between display of bandwidth or IOPs in large fonts on this tile. The symbol means that the number displayed is the average taken during the last *n* seconds. *n* can be configured in **Dashboard I/O Workload** in this window. The symbol means that the number displayed is from the last data sample that was taken. The period between automatic refreshes can be configured in **Dashboard I/O Workload** in this window.
System Clock	Show system clock on the Dashboard
Show Property Sheet in advanced mode	Displays additional details in Property Sheets: • Capacity section—Snapshot Capacity Reserved • Rebuild/Rebalance—Data Movement Jobs • RAM Read Cache—Cache Evictions, Cache Entry, and Cache Skip tables These details are usually only relevant for advanced users and technical support purposes.
Frontend Internals: Show VVols	Displays additional information for VVols
Backend Internals: Show internal Backend views	Displays additional options for Backend table views. These options are recommended only for advanced users and technical support purposes.
Support: Log level	Controls the type of data saved in system logs, which may be required by Customer Support for troubleshooting purposes. The default setting

Table 31 User Preferences (continued)

Item	Description
	recommended for regular operation is Info. Other options include: Trace, Debug, Warn, and Error. **Note** Trace and Debug options may affect system performance, and are usually only recommended for technical support purposes. For more information about logs, see your system's Log Collection Technical Notes.

Modifying system settings

This section explains how to modify system settings after deployment has been performed, including management and data network IP addresses, ESRS, SNMP, and licensing.

Note

For more information about changing security and vCenter configurations, refer to the *VxFlex OS Deployment Guide*.

Note

For more information about security-related settings, refer to the Security chapter in this guide.

Management and Data Network IP addresses

After a VxFlex OS system has been deployed, you can use the VxFlex OS GUI to add management or data IP addresses to the system.

Procedure

1. In the VxFlex OS GUI, from the **System Settings** menu at the top right side of the window, select the **System Settings** option.

   The **System Settings** window is displayed.

2. Select the type of network to which you want to allocate IP addresses: **Management Network** or **Data Network**.

   - The AMS uses management networks to communicate with and control the appliance objects. No I/O passes through these networks. Management networks can be on different subnets, as long as the AMS can communicate with them either via local IP addresses or using a router. If a router is used, the discovery server must be on the same Level 2 physical network as the appliance nodes in order to enable 2-layer broadcasts.
   - For data networks, allocate IP addresses for data networks A and B, 10/25/100 GbE networks that are used for inter-node communications, as well as for passing I/O. For ESXi servers, define ranges for both the ESXi host and the Storage VM (SVM) that resides on each host.

3. In the appropriate parts of the window, do one of the following:
   a. To change the range of an existing IP address pool, edit the addresses.
   b. To add another IP address pool, click the **+** button, and type the new IP address range in the new row.

      Enter a range of IP addresses that is at least equal to the number of servers being added.
4. When you have finished, click **Save**.

Default DNS and NTP IP addresses

After VxFlex OS has been deployed, you can use the VxFlex OS GUI to add default DNS and NTP IP addresses to the system.

Procedure

1. In the VxFlex OS GUI, from the **System Settings** menu at the top right side of the window, select the **System Settings** option.

   The **System Settings** window is displayed.
2. In the panel on the left, click **DNS and NTP**.
3. Select the type of server to which you want to allocate IP addresses: **DNS Default IP Addresses** or **NTP Default IP Addresses**.
4. In the **DNS Default IP Addresses** and **NTP Default IP Addresses** boxes, enter the desired IP addresses.

   Note

   Only IP addresses can be entered in these boxes. Do not use host names.
5. To apply on ESX servers, ensure that the **Apply on ESX** check boxes are selected.
6. When you have finished, click **Save**.

ESRS

This section explains how to configure connection properties, to allow communication between the VxFlex OS system and ESRS (EMC Secure Remote Support) servers. ESRS can be used by ASP Partners for remote hardware support. For more information about using ESRS, refer to ESRS documentation.

You can only use the ESRS feature if you have a valid VxFlex OS license installed.

Before you begin

Before you begin

Before configuring ESRS, you need the following:

- One or more IP addresses of the ESRS servers
- ESRS user and password credentials
- The AMS IP address to be used as the Connect-In IP address. It must be an IP address that is accessible from the ESRS Gateway (for example, in case of NAT). To configure ESRS settings, perform the following:

Procedure

1. From the **System Settings** menu at the top right side of the window, select the **System Settings** option.

   The **System Settings** window is displayed.

2. In the left panel, click **ESRS**.

   The **Set ESRS Server** panel is displayed on the right side of the window.

3. To configure ESRS properties, do one of the following:

   a. To add the first ESRS server, type its IP address in the **ESRS IP Addresses** box.

   b. To add another ESRS server, click the **Add** button .

      A new row is displayed. Type the additional IP address in the new box.

   c. To remove an ESRS server, click the **Remove** button next to the item that you want to remove.

4. In the **User Name** box, type the ESRS user name.
5. In the **Password** box, type the corresponding password for that user.
6. In the **Connect-In IP Address** box, do one of the following:

   a. To add the IP address, type the AMS's IP address.

   b. To remove an IP address, click the **Remove** button next to the IP address.

7. If you do want to use the ESRS feature, and you want to disable alert messages, select the **Allow configuration with no ESRS registered** check box.
8. Click **Save**.
9. Edit the esrs.properties file and add the SWID to the `ForcedSWID` field.

   - Linux: `/opt/emc/scaleio/ams/webapps/ROOT/WEB-INF/classes/`
   - Windows: `C:\Program Files\EMC\scaleio\AMS\webapps\ROOT\WEB-INF\classes`

10. Restart the AMS service.

Results

ESRS configuration properties have been saved in the VxFlex OS system. To validate the connection, proceed to Validating ConnectEMC Dial-Home on page 240.

Validating ConnectEMC Dial-Home

Verify that ConnectEMC (ESRS) is correctly configured and is communicating with Dell EMC Global Services, by performing a dial-home test.

This procedure describes how to quickly verify the connectivity between the ESRS-GW and ESRS servers, either after installation or after registering to ESRS. The ESRS-GW is located at the customer's site and has a verified connection to ESRS servers as part of its installation.

Procedure

1. While logged in to the ESRS-gateway as the gpadmin operating system user, remove the contents of the directory `/opt/connectemc/archive`.
2. Create a test event for ConnectEMC using the following command:

```
touch /opt/connectemc/poll/testerror
```

Results

Within up to 10 minutes (depending on your polling interval setting), ConnectEMC (ESRS) will generate a file in one of the following directories, depending on the test result:
If the connection was successful, a file will be saved in `/opt/connectemc/archive`.

If the service fails, a file will be saved in `/opt/connectemc/archive/failed`.

SNMP

Configure SNMP features using the VxFlex OS GUI.

The following items can be configured:

- IP address and port number of the SNMP Trap Receiver—the IP addresses of up to two Trap Receivers can be added to the system. The default port numbers used for SNMP traps is 162 by default, but can be modified, if necessary.
- Frequency of resending traps, in minutes—the traps are calculated based on MDM polling. A repeating trap will be sent the first time that an alert occurs, and will only be sent again if the resend interval has passed since it was last sent.

Note

For detailed information about SNMP support in the system, see SNMP Trap Support on page 281.

To configure SNMP properties, perform the following steps:

Procedure

1. From the **System Settings** menu at the top right side of the window, in any view, select the **System Settings** option.

   The **System Settings** window is displayed.
2. In the left panel, click **SNMP**.

   The **Set SNMP Configuration** panel is displayed on the right side of the window.
3. To configure Trap Receivers, do one of the following:
   a. To add the first Trap Receiver, type its IP address in the **Trap Receiver** box, and the port number in the **Port** box.
   b. To add another Trap Receiver, click the **Add** button .

      A new row is displayed. Type the additional IP address in the new **Trap Receiver** box, and the port number in the new **Port** box.

c. To remove a Trap Receiver, click the **Remove** button next to the item that you want to remove.

4. Type a number in the **Resend Frequency** box.
5. Click **Save**.

   The SNMP configuration has been added to the system.
6. If you want to send a test SNMP trap, click **Test**.

Install the license - using the VxFlex OS GUI

Before you begin

Before you begin this procedure, ensure that you have the license key from Dell EMC. Getting a license key involves receipt of an LAC (License Authorization Code) email from Dell EMC, and then either downloading the license key from the Dell EMC website, or receiving it by email from Dell EMC.

For more information about obtaining a license key, see the "Licensing" chapter.

This topic explains how to add the license to an existing VxFlex OS system. A license is mandatory for a VxFlex OS system used in a production environment.
To add a license to the system, perform the following steps:

Procedure

1. From the **System Settings** menu at the top right side of the window, in any view, select the **System Settings** option.

   The **System Settings** window is displayed.
2. In the left panel, click **License**.

   The **Set License** panel is displayed on the right side of the window.
3. From the license activation email, or from the downloaded file, copy the license key and paste it in the **Update License** box.

   a. When copying from the license key file (`*.lic`), open the file in a text editor, and copy all the text in the file.
4. Click **Save**.

Results

The license has been added to your system.

Setting the vCenter

In the **System Settings**, you can configure the vCenter.

When using a Linux system, it is recommended to configure the vCenter when there are ESX SDCs connected to the storage system.

Procedure

1. From the **System Settings** menu at the top right side of the window, in any view, select the **System Settings** option.

   The **System Settings** window is displayed.
2. In the left panel, click **vCenter**.

   The **Set vCenter** panel is displayed on the right side of the window.

3. Enter the vCenter IP address and the credentials.

   You can test the connection by clicking **Test Connection**.
4. Select the Data Center.
5. Click **Save**.

Discovery

This topic explains how to configure the connection properties for an Auto Discovery proxy server. A proxy server is only necessary if the VxFlex OS system's AMS and nodes are located on different subnets. Ensure that the relevant AMS software is installed on each proxy server.

Note

In rare cases where nodes are located on a number of subnets, each subnet must have its own Auto Discovery proxy server. After each Auto Discovery operation, change the proxy server's IP address in this window to the proxy server on the next subnet.

To configure the connection properties for an Auto Discovery proxy server, perform the following steps:

Procedure

1. From the **System Settings** menu at the top right side of the window, in any view, select the **System** option.

   The **System Settings** window is displayed.
2. In the left panel, click **Discovery**.

   The **Set Discovery Configuration** panel is displayed on the right side of the window.
3. In the **Auto Discovery proxy IP address** box, type the proxy server's IP address.
4. Select the **Use Auto Discovery** check box.
5. Click **Save**.

Results

The Auto Discovery proxy server connection properties have been added to your system.

Changing login credentials for existing Management users

Use the VxFlex OS GUI to change login credentials for the host Operating System user, SVM user, and BMC (iDRAC) user.

Before you begin

The host operating system user, SVM user, and BMC (iDRAC) user credentials are configured during initial system deployment. If you need to change any of them, do so using the VxFlex OS GUI.

Procedure

1. From the **System Settings** menu at the top right side of the window, select the **System Settings** option.

   The **System Settings** window is displayed.

2. In the left panel, click **Technical Users**.

   The **Set Credentials for Existing Technical Users** panel is displayed on the right side of the window.

3. Click **Management Users**.

   The credentials currently defined for management users are displayed.

4. For the required management user type, type the new username and password in the corresponding boxes.
5. Click **Save**.

Results

The changes that you made to the credentials are saved in the system.

PART 4

Reference

The chapters in this part of the guide describe various topics related to advanced management.

Chapters include:

Chapter 11, "Common Tasks"

Chapter 12, "System events"

Chapter 13, "SNMP Trap Support"

Chapter 14, "Server Factory Reset and Power Operations"

Chapter 15, "DTK - Hardware Update Bootable ISO"

CHAPTER 11

Common Tasks

The following topics describe common tasks that are performed when working with VxFlex Ready Node.

- Associating VxFlex OS volumes with physical disks....248
- Volume information - Linux.... 248
- Port usage and changing default ports....249
- Upgrading Java....250
- Move ESXi hosts to a new vCenter and reconfigure AMS....250
- Best practices for using VMware vMotion.... 251
- Update the hardware using the iDRAC virtual console.... 253
- AMS management takeover of non-managed nodes.... 255

Associating VxFlex OS volumes with physical disks

This section describes how to associate volumes with physical disks.

Contact VxFlex OS Customer Support for access to the troubleshooting utility.

To get VxFlex OS volume information, run the `scli --query_all_volumes` (or `--query_all` or `--query_volume`) command.

Output similar to the following appears:

```
Query-all-volumes returned 10 volumes
Protection Domain 0728185d00000000 Name: pd1
Storage Pool ad99eaab00000000 Name: default
<No volumes defined>

Storage Pool ad99eaac00000001 Name: sp1
Volume ID: fac22a6300000000 Name: vol0 Size: 152.0 GB (155648 MB) Mapped to 1 SDC Thick-provisioned
Volume ID: fac22a6400000001 Name: vol1 Size: 400.0 GB (409600 MB) Mapped to 1 SDC Thin-provisioned
Volume ID: fac22a6500000002 Name: vol2 Size: 80.0 GB (81920 MB) Mapped to 1 SDC Thick-provisioned
Volume ID: fac22a6600000003 Name: vol3 Size: 392.0 GB (401408 MB) Mapped to 1 SDC Thick-provisioned
Volume ID: fac22a6700000004 Name: vol4 Size: 96.0 GB (98304 MB) Mapped to 1 SDC Thin-provisioned
Volume ID: fac22a6800000005 Name: vol5 Size: 112.0 GB (114688 MB) Mapped to 1 SDC Thick-provisioned
Volume ID: fac22a6900000006 Name: vol6 Size: 96.0 GB (98304 MB) Mapped to 1 SDC Thin-provisioned
Volume ID: fac22a6a00000007 Name: vol7 Size: 176.0 GB (180224 MB) Mapped to 1 SDC Thick-provisioned
Volume ID: fac22a6b00000008 Name: vol8 Size: 272.0 GB (278528 MB) Mapped to 1 SDC Thick-provisioned
Volume ID: fac22a6c00000009 Name: vol9 Size: 360.0 GB (368640 MB) Mapped to 1 SDC Thin-provisioned
```

This output shows the Volume ID and name, as well as other volume information.

Volume information - Linux

On the SDC host, run the following command to get the operating system volume information that correlates to the VxFlex OS scini device name:

```
ls -l /dev/disk/by-id/ |grep scini
```

Output, similar to the following appears:

```
lrwxrwxrwx 1 root root 12 Aug 25 19:40 emc-vol-62c093a52d14aec7-fac22a6300000000 -> ../../scinia
lrwxrwxrwx 1 root root 12 Aug 25 19:40 emc-vol-62c093a52d14aec7-fac22a6400000001 -> ../../scinic
lrwxrwxrwx 1 root root 12 Aug 25 19:40 emc-vol-62c093a52d14aec7-fac22a6500000002 -> ../../scinib
lrwxrwxrwx 1 root root 12 Aug 25 19:41 emc-vol-62c093a52d14aec7-fac22a6600000003 -> ../../scinie
lrwxrwxrwx 1 root root 12 Aug 25 19:41 emc-vol-62c093a52d14aec7-fac22a6700000004 -> ../../scinid
lrwxrwxrwx 1 root root 12 Aug 25 19:42 emc-vol-62c093a52d14aec7-fac22a6800000005 -> ../../scinif
lrwxrwxrwx 1 root root 12 Aug 25 19:42 emc-vol-62c093a52d14aec7-fac22a6900000006 -> ../../scinig
lrwxrwxrwx 1 root root 12 Aug 25 19:42 emc-vol-62c093a52d14aec7-fac22a6a00000007 -> ../../scinii
lrwxrwxrwx 1 root root 12 Aug 25 19:42 emc-vol-62c093a52d14aec7-fac22a6b00000008 -> ../../scinih
lrwxrwxrwx 1 root root 12 Aug 25 19:43 emc-vol-62c093a52d14aec7-fac22a6c00000009 -> ../../scinij
```

This output shows the scini volume name and the volume ID.

By matching the volume ID in both outputs, you can match the operating system names, sciniX, with the VxFlex OS volume name.

For example:

- scinia = fac22a6300000000 = vol0
- scinic = fac22a6400000001 = vol1

Alternatively, run the `sg_inq /dev/sciniX` SCSI query command (requires that the sg3_utils be installed on the Linux host). The result of this command includes the EMC volume ID at the bottom of the output, as illustrated:

```
Vendor identification: EMC
Product identification: ScaleIO
Product revision level: 1.3
Unit serial number: EMC-62c093a52d14aec7-fac22a6300000000
```

Note

The product identification remains as ScaleIO (not VxFlex OS).

Port usage and changing default ports

The following table lists the TCP ports that are used by VxFlex OS. Prior to installing or upgrading a system, ensure that these ports are not in use by other processes.

If they are in use, either free them or change them to another available port.

Table 32 Default ports

Port used by	Port #	Protocol	File to change	Field to modify (or to add, if it does not exist)	Notes
AMS and MDM listener	6611	Protobuf over TCP	**Note** Cannot be modified, and must be available		
MDM Cluster member	9011	Protobuf over TCP	`/opt/emc/scaleio/mdm/cfg/conf.txt`	`actor_cluster_port=`*<NEW_PORT>*	
SDS listener	7072	Proprietary protocol over TCP	`/opt/emc/scaleio/sds/cfg/conf.txt`	`tgt_port=`*<NEW_PORT>*	SDCs connect through this port for data communication and to the MDM for meta-data communication. When multiple SDSs are installed on the same physical server, use ports 7072+*x*, where *x* is the index of the SDS (for example, 70721, 70722).
LIA listener	9099	Protobuf over TCP	`/opt/emc/scaleio/lia/cfg/conf.txt`	`lia_port=`*<NEW_PORT>*	The VxFlex OS Installation Manager connects to the LIA to perform installation and maintenance-related operations.
SNMP	162	SNMP v2 over UDP			SNMP traps for system alerts are sent to a trap receiver via this port. The VxFlex OS Gateway sends messages to: snmp.traps_receiver_ip on the port snmp.port
SDBG for MDM (Manager)	25620				Used by VxFlex OS internal debugging tools to extract live

Table 32 Default ports (continued)

Port used by	Port #	Protocol	File to change	Field to modify (or to add, if it does not exist)	Notes
SDBG for MDM (Tie Breaker)	25600				information from the system for debugging purposes. When multiple SDSs are installed on the same physical server, use ports 2564+*x*, where *x* is the index of the SDS (for example, 25641, 25642).
SDBG for SDS	25640				
AMS (secure)	443		`AMS installation directory\conf\server.xml`	`redirectPort=`*<NEW_PORT>* `Connector port=`*<NEW_PORT>*	After changing the port, you must restart the VxFlex OS Gateway service/daemon: • Linux: Run `service scaleio-gateway restart` • Windows: Restart the VxFlex OS Gateway service
AMS (not secure)	80		`AMS installation directory\conf\server.xml`	`Connector port=`*<NEW_PORT>*	

Upgrading Java

Before changing the Java version of a node that is running the VxFlex OS Gateway or AMS of VxFlex OS v2.5 or later, you must prepare lockbox-related files.

The lockbox in VxFlex OS v2.5 saves files in the Java folder of the VxFlex OS Gateway and the AMS. These files need to be saved before any Java version update, then pasted back into the folder.

Procedure

1. From the `jre\lib\ext` (or `jre/lib/ext` for Windows) Java folder, copy these files to a different folder:
   - `commons-lang3-3.6.jar`
   - `cryptoj-6.2.3.jar`
2. Update the Java version.
3. Paste these files back to the folder from where you copied them.

Move ESXi hosts to a new vCenter and reconfigure AMS

Move ESXi hosts to a different vCenter, and reconfigure the AMS accordingly.

Before you begin

Note

If your system uses the AMS Network and Compute features, Customer Support professionals should perform this procedure.

The following steps summarize the workflow for moving ESXi hosts to a different vCenter.

Procedure

1. Verify that the target vCenter is up and running, and is accessible.
2. Connect the ESXi hosts to the new vCenter.

   Refer to VMware guidelines on how to perform this, based on your vCenter version.

   The AMS will display alerts indicating that hosts are not visible in the vCenter.
3. On the AMS, configure the IP addresses and credentials of the new vCenter. For more information, see "Setting the vCenter" in the *VxFlex Ready Node User Guide*.
4. Verify that the alerts in the AMS are no longer present.

Best practices for using VMware vMotion

To use VMware vMotion, set the vMotion status to Enabled for every ESX.

The following ports can be used for vMotion on the ESX. The recommended priority indicates which ports VxFlex OS recommends be used in order to limit the impact on VxFlex OS performance:

Recommended priority	Port	Role	Performance impact
P1	1G port, right port onboard	Unused	None
P2	Bond option of both 1G ports together	AMS management network	May impact AMS management traffic (though negligibly in most cases)
P3	Client 10/25/100 GbE ports: 10/25/100 GbE right onboard and 10/25/100 GbE right on PCI-E slot	Customer network	May impact customer traffic
P4	1G management port, 1G left port onboard	AMS management network	Minimal impact to AMS management traffic

Figure 47 Network ports

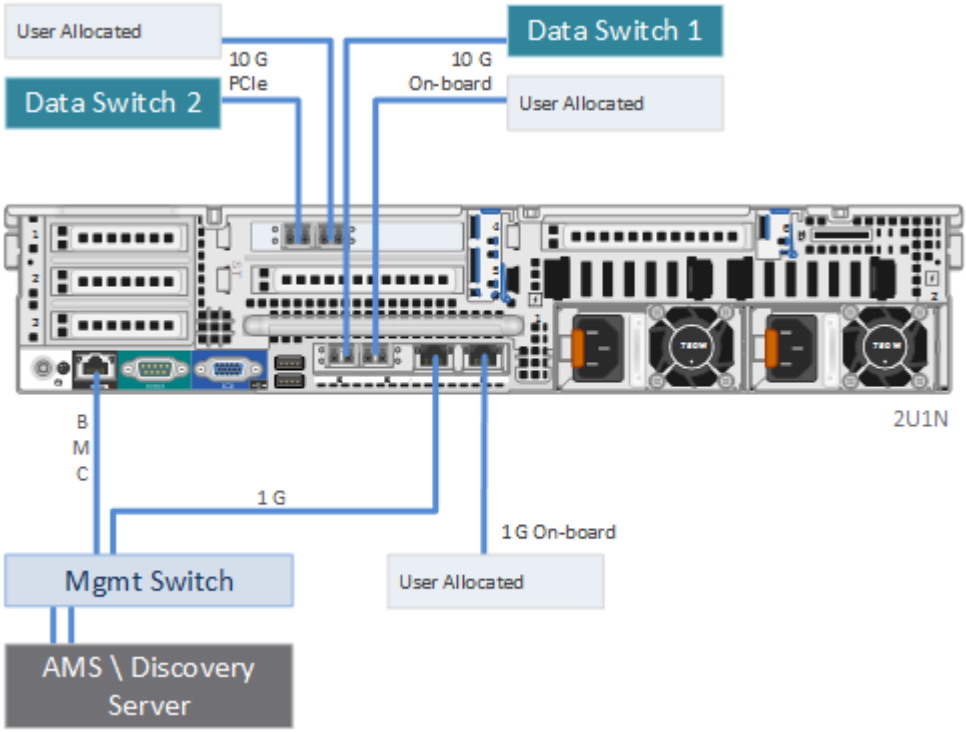

Disable VMware Admission Control

Enable ESXi servers to enter maintenance mode by disabling the VMware Admission Control feature.

Admission Control must be disabled before upgrading a VxFlex Ready Node system. If the ESXi cluster was configured by the VxFlex Ready Node Compute feature with HA enabled, this setting will be configured correctly automatically.
This procedure is relevant for administrators that create an ESXi cluster either not via the AMS orchestration, or via AMS but with HA disabled.

Procedure

1. Disable Admission Control:

Client	Procedure
vSphere Web Client	a. From the vSphere Web Client, log in to the vCenter. b. Select the cluster that contains the VxFlex Ready Nodes. c. Select the **Configure** tab. d. Select **Services** > **vSphere Availability**. e. At the bottom of the window, expand **Admission Control**, and ensure that it is **Disabled**.
vSphere Client	a. From the vSphere client, log in to the vCenter. b. Right-click a cluster that contains VxFlex Ready Nodes and select **Edit Settings**.

Client	Procedure
	c. Click the **vSphere HA** menu, and ensure that **Admission Control** is set to **Disable**.

Update the hardware using the iDRAC virtual console

The iDRAC Virtual KVM console and Virtual CDROM features, provided by the iDRAC Enterprise license, eliminate the need for physical access to the VxFlex Ready Node servers. The Hardware ISO can be attached as a remote/Virtual CDROM image and configured to perform hardware updates to the VxFlex Ready Node firmware.

Note

For additional instructions regarding version updates following SATADOM replacement, see Update the SATADOM firmware (13G servers only) on page 254.

Procedure

1. Make sure that the laptop or server is configured with access to the iDRAC network.

   For details, see "Set up the BMC (iDRAC) IP address and BIOS" in this publication.

2. Download the VxFlex Ready Node Hardware Update Bootable ISO from https://support.emc.com/products/41077https://support.emc.com/products/42216 and make it accessible in the network share folder.
3. Connect to the iDRAC Virtual KVM console.
4. Attach the Hardware ISO to Virtual Media.
   a. From the **Virtual Media** menu, select **Connect Virtual Media**.
   b. From the **Virtual Media** menu, select **Map CD/DVD**.
   c. Click **Browse**, and select the Hardware ISO.
5. From the **Next Boot** menu, select **Virtual CD/DVD/ISO** and configure the Next Boot option so that the server will boot to the iDRAC Virtual CDROM.
6. Power-cycle the server using the appropriate **Power** menu option, and allow the server to boot to the virtual media.
7. Repeat steps 3 on page 253 to 6 on page 253 for each VxFlex Ready Node server that requires BIOS or firmware updates and configuration. All VxFlex Ready Node servers can be updated in parallel.
8. When the updates are completed, refresh the iDRAC browser screen, log in to the iDRAC, and re-launch the virtual console as needed.
9. Wait for the configuration and firmware updates to complete. The server console screen will indicate when the script is complete.

   WARNING

   Do not reboot the VxFlex Ready Node server while the update process is being performed!

   The iDRAC will be reset several times during the update process. This causes the iDRAC virtual console viewer to close, virtual media to disconnect, and the

iDRAC browser window to be unavailable for several minutes during each reset. The hardware update scripts will continue to run from RAM on the server.

The update script will generate a log indicating whether each configuration and firmware flash is successful.

10. (Optional) Check each VxFlex Ready Node server's log for successful competition:
    a. After the update script completes, press Alt+F2 to access a user console, and then Enter to log in.
    b. Check the log contents for errors:

    ```
    less /bundleapplicationlogs/apply_components.log
    ```

    For more information, see Troubleshooting the Hardware ISO on page 331 .

    c. Press q to exit the log viewer, and then Alt+F1 to access the original console screen.
11. Reboot the VxFlex Ready Node server and allow the update and configuration jobs to complete. Power-cycle the server using the appropriate **Power** menu option.
12. For each VxFlex Ready Node server, after the updates are finalized, clear the iDRAC job queue using the iDRAC GUI:
    a. From your Internet browser, go to https://*<iDRAC_IP_address>*.
    b. In the **DELL Console Login** window, type these credentials:
    - username: root
    - password: *<password>* (for Dell-supplied nodes) or password (for EMC-supplied nodes)

    c. Click **Login**.
    d. In the iDRAC GUI navigation pane, select the **Maintenance** tab, and then select the **Job Queue** tab.
    e. Ensure that all jobs have completed successfully. Any job failures may require re-running the bootable ISO, or further troubleshooting.
    f. Select the checkbox for all of items in the **Job Queue** list, and then click **Delete**.
13. Continue with OS installation and configuration, VxFlex OS deployment, and other required tasks.

Update the SATADOM firmware (13G servers only)

The Dell factory should provide the SATADOM boot device with the correct firmware already loaded. However, in some cases, it may be necessary to update the firmware loaded on the SATADOM boot device.

For convenience, SATADOM firmware version S130710K is included on the VxFlex Ready Node Hardware ISO.

WARNING

The SATADOM firmware update process permanently erases all data from the SATADOM device. Proceed with caution!

Procedure

1. Connect to the iDRAC Virtual KVM console.

   For instructions, see Open the KVM console - 13G servers on page 318.

2. Boot the VxFlex Ready Node server to the Hardware ISO.

   For instructions, see Update the hardware using the iDRAC virtual console on page 253.

3. After the update script completes, press Alt-F2 to access a user console.
4. Press Enter to login.
5. Execute the SATADOM firmware:

```
/tmp/S130710K_read_log_issue/satadom.sh
```

6. Press y to begin the SATADOM update.
7. Press Alt+F1 to access the original console screen, and then press Enter to reboot the server.

AMS management takeover of non-managed nodes

The AMS can "take ownership" of all the nodes on an existing VxFlex Ready Node system that is currently running VxFlex OS without AMS management. When the process is complete, all the servers in the system are managed with the AMS.

The take-ownership process is performed from the VxFlex OS GUI together with an AMS CLI command.

AMS management takeover compatibility

Before starting the take-ownership process, ensure that the non-AMS-managed VxFlex Ready Node system is compatible.

Supported configurations

- Hardware
  - Only VxFlex Ready Node PowerEdge 13G and 14G servers are supported.

    Note

    Only configurations supported by the AMS are supported by this feature.
- Operating systems
  - ESX 6.0.x or ESX 6.5.x
  - RHEL 7.3 or RHEL 7.4
- VxFlex OS
  - Ensure that all nodes are running the same version of VxFlex OS software as the AMS.

- Networking
  - In two-layer, bare metal systems, layer 1 contains SDS, MDM and LIA components, and can use a 2-channel or 4-channel data network. Layer 2 (both in ESXi and Linux systems) consists of SDCs only, on a 2-channel data network.
  - The iDRAC component of each server should be configured.
  - vCenter (when used) must be configured.
- Devices
  - On 14G servers running ESX, NVMe SSD storage devices must be configured as RDM. HDD and other SSD devices must be configured as DirectPath.
  - On 13G servers running ESX, all devices must be configured as DirectPath.

Prerequisites

- On RHEL nodes, ensure that the following packages are installed on every node:
  - libx86emu
  - hwinfo
  - perccli
  - ipmitool
  - smartmontools
  - lshw
  - sg3_utils
  - ntp
  - pciutils
  - sas3ircu (copy the file to an executable path, and assign it execute permissions)
- Each VxFlex OS data network belongs to a unique subnet.
- A VxFlex OS MDM virtual IP address must be configured on the two data networks used by the SDCs.
- All nodes (including SVM) nodes must be accessible via SSH.

Limitations

- Nodes that were previously not successfully installed or properly configured will be ignored during the take ownership process.
- If the current node credentials do not match the VxFlex OS password policy, prepare a CSV file with those credentials. The CSV will be used during the take-ownership procedure. For more information about creating the CSV file, see Creating a management takeover CSV file on page 262.

AMS management takeover of non-managed nodes

Run the AMS take-ownership process, using the VxFlex OS GUI and the AMS CLI.

Before you begin

Ensure that:

- The servers for the VxFlex OS GUI and the VxFlex Ready Node AMS meet the system requirements described in "System requirements."

- The non-AMS-managed system is compatible. For more information, see AMS management takeover compatibility on page 255.
- You know the IP ranges for the existing:
  - Management IP addresses for all nodes and storage VM (SVM, if applicable)
  - Data IP addresses for all nodes and SVM (if applicable)
  - iDRAC IP addresses
- You know the credentials for the following:
  - MDM admin user (or LDAP admin user)
  - iDRAC
  - Storage virtual machine (SVM, if applicable)
  - Host (ESXi/Linux)
  - Data IP addresses
  - Monitor user (if using LDAP)
- All nodes must be reachable.
- In the AMS setup, the default is to use the same credentials (AMS, management, SVM, and iDRAC) for each node. To assign different credentials per node/SVM, prepare a CSV, as described in AMS management takeover compatibility on page 255.

To run the take-ownership process, you must do the following:

1. Install the VxFlex OS GUI.
2. Install the VxFlex Ready Node AMS.
3. Use the VxFlex OS GUI to perform the AMS initial setup.
4. Run and AMS CLI command to perform the take-ownership process.

Procedure

1. Install the VxFlex OS GUI:

   ```
   EMC-ScaleIO-gui-2.6-<build>.X.msi
   ```

   For more information, see "Install the VxFlex OS GUI."

2. Install the VxFlex Ready Node AMS:
   - Windows:

     ```
     EMC-ScaleIO-ams-2.6-<build>-x64.msi
     ```

     To install as a proxy discovery server, select that option during the installation wizard.
   - Linux:

     ```
     rpm -ivh EMC-ScaleIO-ams-2.6-<build>.x86_64.rpm
     ```

     To install as a proxy discovery server, add `INSTALL_ADS=true` before the install command.

When using Linux AMS, restart the AMS server. For more information, see "Install the VxFlex Ready Node AMS."

3. Log in to the VxFlex OS GUI:

   a. Open the VxFlex OS GUI by selecting **Start** > **All Programs** > **VxFlex OS GUI**.

      The initial login screen is displayed.

   b. Type the AMS IP address or hostname and click **Connect**.

      Note

      Ensure that you type the IP address of the AMS server, not of the MDM server.

      If a certificate notice is displayed, review and accept the certificate.

   c. In the login screen, type the AMS username and a temporary password, then click **Login**.

      You are prompted to change the temporary password. Type and confirm a new password, then click **Login**.

   The AMS Initial Setup Wizard begins. For a full description of this process, see "Deploying VxFlex Ready Node."

4. In the **System** screen, set general system preferences:

   a. Type a system name (optional).

   b. Select the operating system environment.

      Your choice here will customize the AMS setup and deployment process for the selected operating system only. You can add nodes of a different operating system type after completing the initial deployment.

      Note

      If the system to be taken over has both Linux and ESXi nodes, select **ESXi** and enter the credentials and IP addresses for those nodes only. Add the Linux nodes after the initial setup, as described at the end of this process.

      This documentation describes steps for all operating systems, even though some steps (such as SVM configuration) are not relevant for Linux deployment.

   c. Select to set up the AMS as a discovery proxy (optional), and enter its IP address.

   d. Click **Next**.

5. In the **Security** screen, type the passwords of the non-AMS-managed system nodes.

   Note

   If you prepared a CSV file for your credentials, the password information in that file will overwrite what is entered in this step. Type passwords for these fields that meet the VxFlex OS criteria.

a. Type the host operating system password.

b. For ESXi systems, type the SVM password.

c. Type the BMC (iDRAC) password.

d. Type a new AMS recovery key.

This password is required for running the AMS recovery process. It must meet the following criteria:

- At least 7 characters long
- Include at least 3 of the following groups: [a-z], [A-Z], [0-9], special chars (!@#$ ...)
- An uppercase character at the beginning of the password or a number at the end of the password is not included in the count of character groups

e. Click **Next**.

6. In the **Management Networks** screen, type IP address ranges for the management networks.

a. Select the IP address version (IPv4 or IPv6) in **IP address version**.

b. Type IP address ranges for each management network, using the options described below.

Note

If the non-AMS-managed system has both ESXi and Linux nodes, only type the addresses for the ESXi hosts in this step.

c. Click **Next**.

To do this	Use this
Use VLAN tagging	Select **VLAN ID** and type the ID. Use VLAN tagging to separate the networks on different VLANs and assign a VLAN ID. The host and the SVM should have the same VLAN ID. Entering a VLAN ID of 0, or leaving the VLAN ID field empty, will not enable VLAN tagging for that network.
Define non-contiguous IP address ranges	Click . Add the range in the line that appears.
Delete a line	Click .

7. In the **Data Networks** screen, type IP address ranges for each data network.

a. Select the IP address version (IPv4 or IPv6) in **IP address version**.

IPv6 is supported in RHEL environments only.

b. To use all available NICs for storage (Linux systems only), select that option.

All available NICs are displayed for configuring.

c. Type the IP address for the cluster virtual IP.

d. Type IP address ranges for the data network, using the same options as those described for the management networks.

Note

If the non-AMS-managed system has both ESXi and Linux nodes, only type the addresses for the ESXi hosts in this step.

e. Type ranges for both the ESX host as well as the SVM that resides on each host.

f. Click **Next**.

8. (Optional) In the **DNS and NTP** screen, type the IP addresses of the DNS and NTP servers, then click **Next**.

Note

Type the IP address, not the hostname.

In RHEL environments, click **Finish**, then the **Add Nodes** option is displayed.

9. (Optional) In the **Virtualization** screen, define the vCenter:

a. Enter the vCenter IP address and credentials.

b. Click **Test connection**.

c. Select a data center to use for the ESX nodes.

d. Click **Finish**.

The AMS setup is complete.

10. If the non-AMS-managed system has both ESXi and Linux nodes, add the Linux nodes (otherwise, skip to the next step):

a. From the VxFlex OS GUI, select **System Settings** > **Management and Data IPs**.

b. Select **Disable range validation (Mixed mode)**.

c. Select the IP address version (IPv4 or IPv6) in **IP address version**.

d. On the **Management Network** screen, type the IP address ranges for the management networks.

e. On the **Data Network** screen, type the IP address ranges for the data networks.

f. Click **Save**.

The Linux nodes from the non-AMS-managed system are ready for AMS takeover.

11. From a command line on the AMS server, run this AMS CLI command:

Operating System of AMS server	Command to run
Windows	`C:\Program Files\EMC\scaleio\AMS\webapps\ROOT\resources\scaleio_repository java -jar EMC-ScaleIO-ams-cli-<VERSION>.jar --username admin --password <ADMIN_PASSWORD> --take_ownership --mdm_ips <MDM_IP_ADDRESSES> --mdm_username admin --mdm_password <MDM_PASSWORD> --monitor_username <MONITOR_USER_NAME> --monitor_password <MONITOR_PASSWORD>`
Linux	`/opt/emc/scaleio/ams/webapps/ROOT/resources/scaleio_repository/scaleio_repository java -jar EMC-ScaleIO-ams-cli-<VERSION>.jar --username admin --password <ADMIN_PASSWORD> --take_ownership --mdm_ips <MDM_IP_ADDRESSES> --mdm_username admin --mdm_password <MDM_PASSWORD> --monitor_username <MONITOR_USER_NAME> --monitor_password <MONITOR_PASSWORD>`

If the non-AMS-managed system is not using LDAP, enter any strings for `--monitor_username` and `--monitor_password`.

For more information on the AMS CLI, see the *VxFlex Ready Node CLI Reference Guide*.

To supply credentials in a CSV file, add the following to the end of the command: `--custom_node_credentials_file "<PATH_TO_CSV_FILE>"`

For example:

```
C:\Program Files\EMC\scaleio\AMS\webapps\ROOT\resources
\scaleio_repository --username admin --password Scaleio123 --
take_ownership --mdm_ips 10.234.215.223 --mdm_username admin
--mdm_password Scaleio123 --monitor_username remander --
monitor_password reman123 --custom_node_credentials_file "c:/
customCredsPlugin2.csv"
```

This process can take several minutes. Progress is displayed on the screen.

12. Replace the certificates of the nodes that you added, using one of the following methods:
    - AMS CLI command: `--replace_nodes_certificates`
    - From the VxFlex OS GUI, select **System Settings** > **Renew Certificates**.

Results

The AMS can now be used to manage all the nodes in the previously non-managed system.

After you finish

Perform the mandatory (and optional) post-deployment steps, described in the "Post-deployment" section.

Creating a management takeover CSV file

During the management takeover process, you can upload a custom CSV file to assign different credentials for each iDRAC or node (RHEL or ESXi).

You can use a custom CSV file in the following circumstances:

- The credentials of nodes to be taken over do not match the VxFlex OS policy. For example, the node password doesn't meet the VxFlex OS requirements.
- The nodes don't all have the same password.

In both cases, in the AMS deployment wizard, you must type a valid password (which is used for future system expansion). In the custom CSV, type the IP address, username, and password of each node. The username and password in the CSV will overwrite those that are assigned by the AMS wizard.

The CSV must contain the following data:

IP address	Username	Password
IP address Management IP address of: • iDRAC • Node (RHEL) • SVM (ESXi) • ESXi	Username for given IP address	Password for given IP address

Example CSV file:

```
10.234.210.2,admin,Scaleio1234!
10.234.210.3,root,Password2
10.234.210.4,administrator,Scaleio1234!
10.234.211.2,root,SScaleio123
10.234.211.3,root,SScaleio124
10.234.211.4,user,password
10.234.209.2,admin,password
10.234.209.3,root,password1
10.234.209.4,root,password2
```

Leading and trailing blank spaces are not supported in the username or password fields. They will be truncated.

CHAPTER 12

System events

The following topics describe VxFlex Ready Node system events and alerts.

- System events overview 264
- Event format 264
- Viewing events 265
- Event list 268

System events overview

This appendix describes event messages that can be generated by VxFlex OS.

An event message is generated as a response to changes that have occurred while the system is running. Event messages to notify you of the changes in case your intervention is needed. Each event message is associated with a severity level. The severity indicates the risk (if any) to the system, in relation to the changes that generated the event message. The severity levels are as follows:

Severity level	Explanation	Example
Info	Informs you of events that one should be aware of, but that do not put the system at risk (no urgency).	`CLI_COMMAND_RECEIVED`
Warning	Indicates a failure that may result from an acceptable condition (e.g. user error), but can also indicate a possible failure.	`SDC_DISCONNECTED`
Error	An error alarm was raised by the system. This error requires your attention. The system is stable, but could be degraded.	`MDM_DATA_DEGRADED`
Critical	A major error alarm was raised by system. The system requires immediate attention. Some data is unavailable.	`MDM_DATA_FAILED`

Event entries are documented as follows:

Name
The name associated with the event

Message
The message that will appear

Severity
The severity level

Description
A description of the reasons that triggered the event notification

Action
Possible actions that can resolve the reported event (if relevant)

Event format

All event messages received are in a parsable structured format, containing the following fields:

ID: a sequential number attached to all events

Date: the local time set in the server

Format: YYYY-MM-DD hh:mm:ss.ssssss

Name: the unique name of the event

Severity: one of the predefined severity levels

Message: message describing the event. Some event notification message verbosity may be expanded by using the full switch (see "Viewing events locally").

The following is an example of a possible event notification:

```
139  2013-07-22 17:21:11.694571 CLI_COMMAND_RECEIVED
INFO     Command MAP_VOL_TO_SCSI_INITIATOR Received
```

This event has the option of extended verbosity. When requested, the event notification will be displayed as follows:

```
139  2013-07-22 17:21:11.694571 CLI_COMMAND_RECEIVED
INFO     Command MAP_VOL_TO_SCSI_INITIATOR Received Vol
Name: snap_raw; SCSI Initiator Name: ini-21
bAllocateLunNum: 1 lunNum: 0
```

The following is a breakdown of the event according to the fields in the event record (as described above):

Parameter	Description
ID	139
Date	2013-07-22 17:21:11.694571
Name	CLI_COMMAND_RECEIVED
Severity	INFO
Message	Command MAP_VOL_TO_SCSI_INITIATOR Received
Extended	Command MAP_VOL_TO_SCSI_INITIATOR Received Vol Name: snap_raw; iSCSI Initiator Name: ini-21 bAllocateLunNum: 1 lunNum: 0

Viewing events

You can view events in the following ways:

- on a local server
- via Syslog
- via email

To configure events via email, see the *EMC Secure Remote Services Installation and Operations Guide*.

Viewing events locally

Events can be viewed by running the following command, and by using switches to filter the data.

Command:

showevents.py

Syntax

```
/opt/emc/scaleio/mdm/bin/showevents.py [Options]
```

Actual command syntax is operating-system dependent. For more information, see CLI basics on page 100.

Description/Notes

Displays events, which can be filtered by optional switches.

Parameters

Parameter	Description
Options:	
--min_severity <SEVERITY>	Displays events with at least the specified minimum severity
--severity <SEVERITY>	Displays events with the specified severity
--from_id <ID>	Displays all events starting from the given ID
--to_id <ID>	Displays all events ending at the given ID
--from_date <ID>	Displays all events starting from the given date
--to_date <ID>	Displays all events ending at the given date
--grep <TEXT>	Displays events containing the specified text
--full	Extends message verbosity

Example

```
/opt/emc/scaleio/mdm/bin/showevents.py --severity ERROR
```

Viewing events in Syslog

The MDM syslog service can send VxFlex OS events, via TCP/IP, to RFC 6587-compliant remote (or local) Syslog servers. Messages are sent with facility `local0`, by default.

Once the syslog service is started, all events will be sent until the service is stopped.

This section describes how to use the CLI to start, stop, and configure the facility field of the syslog events.

Start posting events to remote syslog servers

Command:

`start_remote_syslog`

Syntax

```
scli --start_remote_syslog --remote_syslog_server_ip
<IP> [--remote_syslog_server_port <PORT>]
[--syslog_facility <FACILITY>] [--attach_event_code]
```

Actual command syntax is operating-system dependent. For more information, see "CLI basics".

Description/Notes

Starts posting events to one or more remote syslog servers.

Parameters

Parameter	Description
--remote_syslog_server_ip <IP>	A comma-separated list of syslog server IP addresses (maximum of two servers)
--remote_syslog_server_port <PORT>	The syslog server port (default 1468)
--syslog_facility <FACILITY>	Controls the facility field of the event (legal values 0—23; default 16)
--attach_event_code	Add the posted event code to the event message (disabled, by default)

Example

```
scli --start_remote_syslog --remote_syslog_server_ip 192.168.1.10 --
remote_syslog_server_port 1500
```

Stop posting events to remote syslog servers

Command:

stop_remote_syslog

Syntax

```
scli --stop_remote_syslog --remote_syslog_server_ip
```

<IP>

Actual command syntax is operating-system dependent. For more information, see "CLI basics".

Description/Notes

Stops posting events to remote syslog servers.

Parameters

Parameter	Description
--remote_syslog_server_ip <IP>	A comma-separated list of syslog servers IP addresses

Example

```
scli --stop_remote_syslog --remote_syslog_server_ip
192.168.1.10,192.168.1.20
```

Configure the syslog events facility field

Command:

`set_syslog_facility`

Syntax

```
scli --set_syslog_facility --remote_syslog_server_ip <IP> --
syslog_facility <FACILITY>
```

Actual command syntax is operating-system dependent. For more information, see "CLI basics".

Description/Notes
Configures the facility field of the syslog events.

Parameters

Parameter	Description
--remote_syslog_server_ip <IP>	A comma-separated list of syslog server IP addresses
--syslog_facility <FACILITY>	Controls the facility field of the event (legal values 0—23; default 16)

Example

```
scli --set_syslog_facility --remote_syslog_server_ip
192.168.1.10,192.168.1.20 --syslog_facility 20
```

Event list

This section lists all VxFlex Ready Node events, grouped by the following categories:

- Authentication
- CLI commands
- License and installation
- MDM
- SDC
- SDS
- Rebuild

Authentication

Authentication Failed

Parameter	Description
Name	AUTHENTICATION_FAILED
Message	Authentication failed for user U
Severity	Warning
Description	User entered the wrong password
Action	If you see this event multiple times, someone may be trying to gain unauthorized access to the system.

CLI commands

CLI Command Received

Parameter	Description
Name	CLI_COMMAND_RECEIVED
Message	Command X Received
Severity	Info
Description	CLI command X was entered by a user
Action	None

CLI Command Succeeded

Parameter	Description
Name	CLI_COMMAND_ SUCCEEDED
Message	Command X ended successfully
Severity	Info
Description	CLI command X was executed successfully
Action	None

CLI Command Failed

Parameter	Description
Name	CLI_COMMAND_FAILED
Message	Command X failed with error E
Severity	Warning

Parameter	Description
Description	The CLI command X entered by user failed with error E
Action	Look up error E and address the issue accordingly

Snapshot volumes could not be found, by ID

Parameter	Description
Name	SNAPSHOT_VOLUMES_FAILED_BY_ID
Message	Could not snapshot volumes, because a volume ID was not found. ID: "MOS_OBJID__FORMAT"."
Severity	Error
Description	This message is posted if a snapshot_volume command contains an invalid volume ID (out of many). The CLI will only get an error code, but in the event, you can see which volume ID is invalid.
Action	Verify the parameters entered for the snapshot_volume command

Snapshot volumes could not be found, by name

Parameter	Description
Name	SNAPSHOT_VOLUMES_FAILED_BY_NAME
Message	Could not snapshot volumes, because a volume was not found: Name: %s.
Severity	Error
Description	This message is posted if a snapshot_volume command contains an invalid volume name (out of many). The CLI will only get an error code, but in the event, you can see which volume name is invalid.
Action	Verify the parameters entered for the snapshot_volume command

License and installation

License Expiration Warning

Parameter	Description
Name	LICENSE_EXPIRATION_WARNING
Message	License will expire in X days
Severity	Warning
Description	System license will expire in 30 days or less
Action	Contact EMC Support for license renewal, and then reinstall.

License Expiration Error

Parameter	Description
Name	LICENSE_EXPIRATION_ERROR
Message	License will expire in X days
Severity	Error
Description	System license will expire in 7 days or less
Action	Contact EMC Support for license renewal. If you have already renewed your license, install it.

License Expiration Critical

Parameter	Description
Name	LICENSE_EXPIRATION_CRITICAL
Message	License will expire in X days
Severity	Critical
Description	System license will expire in 2 days or less
Action	Contact EMC Support for license renewal. If you have already renewed your license, install it.

License Expired

Parameter	Description
Name	LICENSE_EXPIRED
Message	License has expired
Severity	Critical

Parameter	Description
Description	The system's license has expired
Action	To resume operational mode, contact EMC Support for license renewal. If you have already renewed your license, install it.

Upgrade has started

Parameter	Description
Name	UPGRADE_STARTED
Message	Upgrade to version %s has started.
Severity	Info
Description	An upgrade procedure has been initiated
Action	Not needed

Upgrade has finished

Parameter	Description
Name	UPGRADE_FINISHED
Message	Upgrade completed successfully.
Severity	Info
Description	An upgrade procedure completed successfully
Action	Not needed

Upgrade has failed

Parameter	Description
Name	UPGRADE_FAILED
Message	Upgrade was not successful. Reason: %s
Severity	Error
Description	An upgrade procedure was not able to complete
Action	Fix the error and retry the upgrade

MDM

MDM Started

Parameter	Description
Name	MDM_STARTED
Message	MDM Process started
Severity	Info
Description	MDM process has started running
Action	None

MDM Data Degraded

Parameter	Description
Name	MDM_DATA_DEGRADED
Message	Some of the Storage Pool data is now in Degraded state
Severity	Error
Description	Some of the Storage Pool data is in Degraded state. This data is not protected against another failure.
Action	The system is rebuilding the Degraded data to return to Normal (protected) state. Check if any hardware is malfunctioning and requires replacement.

MDM Data Failed

Parameter	Description
Name	MDM_DATA_FAILED
Message	Some Storage Pool data is now unavailable
Severity	Critical
Description	Multiple failures have occurred. Some Storage Pool data is now unavailable. This data cannot be accessed.
Action	Locate and fix the failed hardware. If the problem is not resolved, contact EMC Support.

MDM Data Normal

Parameter	Description
Name	MDM_DATA_NORMAL
Message	All of the Storage Pool data has returned to Normal state
Severity	Info
Description	All Storage Pool data previously in Degraded or Failed state has returned back to Normal state. User data is fully accessible and protected.
Action	None

SDC

New SDC Connected

Parameter	Description
Name	NEW_SDC_CONNECTED
Message	New SDC (IP: X; ID: Y; GUID: Z) connected
Severity	Warning
Description	A new SDC (IP: X; ID: Y; GUID: Z) has connected to the MDM
Action	A new SDC has just connected to the MDM. Validate that this is a valid SDC.

SDC Connected

Parameter	Description
Name	SDC_CONNECTED
Message	SDC (IP: X; ID: Y; GUID: Z) reconnected
Severity	Info
Description	An existing SDC (IP: X; ID: Y; GUID: Z) has reconnected to the MDM
Action	None

SDC Disconnected

Parameter	Description
Name	SDC_DISCONNECTED
Message	SDC (IP: X; ID: Y; GUID: Z) disconnected
Severity	Warning
Description	SDC (IP: X; ID: Y; GUID: Z) has disconnected from the MDM.
Action	Make sure it is expected otherwise this might be a hardware malfunction

SDS

SDS Disconnected

Parameter	Description
Name	SDS_DISCONNECTED
Message	SDS X (IP:Y; ID: Z) disconnected
Severity	Error
Description	SDS X (IP: Y; ID Z) has disconnected from the MDM
Action	Make sure that this is an expected event, because otherwise this might be caused by a hardware malfunction.

SDS Reconnected

Parameter	Description
Name	SDS_RECONNECTED
Message	SDS X (IP:Y; ID: Z) reconnected
Severity	Info
Description	SDS X (IP: Y; ID Z) has reconnected to the MDM. If this event appears multiple times subsequently for the same SDS (and not directly after SDS_DISCONNECTED), it can indicate a bad network connection.
Action:	Check network connections.

SDS Remove Done

Parameter	Description
Name	SDS_REMOVE_DONE
Message	SDS X (IP:Y; ID: Z) was removed successfully
Severity	Info
Description	The asynchronous process of removing an SDS has completed.
Action	None

Open SDS Device Failed

Parameter	Description
Name	OPEN_SDS_DEVICE_FAILED
Message	Failed to open a device D on SDS X (IP:Y; ID: Z) with error message E
Severity	Error
Description	Failed to open storage device D on SDS X (IP:Y; ID: Z) with error message E
Action	Check the cause of the error, and identify if it's a human error or a system malfunction. Check hardware if needed.

SDS Device Error Report

Parameter	Description
Name	SDS_DEV_ERROR_REPORT
Message	Device error reported device D on SDS X (IP:Y; ID: Z)
Severity	Error
Description	Error reported on storage device D on SDS X (IP:Y; ID: Z)
Action	Check the storage device on the server

Device capacity is high

Parameter	Description
Name	DEV_CAPACITY_USAGE_HIGH
Message	Capacity usage on %s is HIGH.")

Parameter	Description
Severity	Warning
Description	Capacity is high, due to capacity used by snapshots/thin volumes
Action	Remove unnecessary snapshots or add more storage

Device capacity is critical

Parameter	Description
Name	DEV_CAPACITY_USAGE_CRITICAL
Message	Capacity usage on %s is CRITICAL.")
Severity	Error
Description	Capacity is critical, due to capacity used by snapshots/thin volumes
Action	Remove unnecessary snapshots or add more storage

Device capacity has returned to normal

Parameter	Description
Name	DEV_CAPACITY_USAGE_NORMAL
Message	Capacity usage on %s is normal.")
Severity	Info
Description	Capacity usage is back to normal
Action	Not needed

SDS configuration has become invalid

Parameter	Description
Name	SDS_CONFIG_INVALID
Message	CLI_TARGET_NAME_CAP" %s (ID "MOS_OBJID__FORMAT ") configuration is invalid.
Severity	Critical
Description	The SDS cannot access its configuration files.
Action	Contact EMC Support

SDS disk errors were fixed

Parameter	Description
Name	SDS_FIX_DISK_ERROR
Message	CLI_TARGET_NAME_CAP" %s (ID "MOS_OBJID__FORMAT ") device %s fixed %d disk errors via reads.
Severity	Warning
Description	There were read errors on this device that were fixed by reading the data from secondary and re-writing it. This may be a sign of impending device hardware malfunction.
Action	Check for hardware malfunction

Background device scanner comparison error

Parameter	Description
Name	SCANNER_COMPARE_REPORT
Message	Background device scanner on device ID 2301536800030001 reported compare error (Device Path: <device path>, SDS: <SDS name and ID>, Peer Device Path: <peer device path>, Peer SDS: <peer SDS name and ID>, Volume name: <volume name>, Volume offset: <volume offset>)
Severity	Error
Description	Background device scanner error report, which provides details about comparison errors found during comparison of two copies of data on different devices.
Action	Check storage device for hardware malfunction

Rebuild

No Rebuild Progress Warning

Parameter	Description
Name	NO_REBUILD_PROGRESS _WARNING
Message	No rebuild progress for 30 minutes
Severity	Warning

Parameter	Description
Description	Rebuild did not progress for 30 minutes during the current recovery
Action	Contact EMC Support

No Rebuild Progress Error

Parameter	Description
Name	NO_REBUILD_PROGRESS _ERROR
Message	No rebuild progress for 60 minutes
Severity	Error
Description	Rebuild did not progress for 60 minutes during the current recovery
Action	Contact EMC Support

No Rebuild Progress Critical

Parameter	Description
Name	NO_REBUILD_PROGRESS _CRITICAL
Message	No rebuild progress for 180 minutes
Severity	Critical
Description	Rebuild did not progress for 180 minutes during the current recovery
Action	Contact EMC Support

Rebuild Progress Resumed

Parameter	Description
Name	REBUILD_PROGRESS_RESUMED
Message	Rebuild progress resumed
Severity	Info
Description	Following a detection of a rebuild not progressing, the system has now detected that the rebuild progress has resumed. The system is currently recovering.
Action	None

CHAPTER 13

SNMP Trap Support

The following topics describe VxFlex Ready Node support for SNMP.

- General 282
- Supported alerts and event numbering conventions 282
- Configure SNMP properties after deployment 308
- VxFlex OS.mib file 310

General

SNMP traps are implemented as part of VxFlex OS, using SNMP v2. UDP transport is used for SNMP, and the default port for trap communication is 162. The SNMP feature is enabled by default.

The SNMP trap sender includes a proprietary/custom `MIB` called `scaleio.mib`. This `MIB` file is located on the VxFlex OS gateway server, in the `/gateway/webapps/ROOT/WEB-INF/classes` folder. A copy of the `MIB` file is included in at the end of this section. A general trap type with a unique identification number (OID) is defined in the `MIB`, so that the SNMP traps are configured to contain alert data within themselves, and use a single OID (as opposed to granular traps). All the SNMP traps contain variable bindings for severity type, which is the alert classification; the ID of the source object for which the alert was created; and an action code, which is the event number.

Only TRAP commands/messages are supported, and are initiated by the VxFlex OS SNMP traps manager. GET/SET operations are not supported (or more specifically, GET/GET NEXT/GET BULK/SET/INFORM/RESPONSE).

In addition to SNMP traps, alert messages are also displayed in the VxFlex OS GUI.

Supported alerts and event numbering conventions

The following alerts can be sent as SNMP traps by VxFlex OS. All events are numbered in the following format: SIO<CLASS>.<TYPE>.<ISSUE>. The issue number is a running counter for all issues in a specific type.

Open and closing alerts will consist of the same code and issue number, with the exception of the first digit (0 or 1) in the <ISSUE> section. For example:

- SIOXX.XX.0XXXXXX indicates that the alert is active
- SIOXX.XX.1XXXXXX indicates that the alert has been closed

CLASS/TYPE:

System = 1

- Capacity = 1
- License = 2

MDM = 2

- MDM_Cluster = 1
- Protection_Domain = 2
- Fault_Set = 3
- Storage_Pool = 4

SDS = 3

- SDS = 1
- Device = 2

SDC = 4

- SDC = 1

Volume = 5

Storage Controller = 6

- Physical Drive = 1
- Logical Disk = 2
- Battery = 3

Node = 7

- Boot Drive = 1
- Storage Controller = 2
- CPU Socket = 3
- RAM = 4
- Node = 5

CPU Socket = 8

- Socket Cache = 1

Chassis = 9

- Fan = 1
- PSU = 2

ESRS = 10

- ESRS = 1

NIC= 11

- NIC Port = 1

The same conventions will be applied in the future to additional system events.

The following table describes each alert, and the actions to be taken for each one:

Note

Each alert has a corresponding Closed state, represented by the code SIOXX.XX.1XXXXXX. For example Open state: SIO02.01.0000003, Closed state: SIO02.01.1000003

VxFlex OS Alerts in SNMP, VxFlex OS GUI, REST, and ESRS

This table summarizes alerts generated by VxFlex OS systems.

Table 33 VxFlex OS Alerts in SNMP, VxFlex OS GUI, REST, and ESRS

Alert Message in VxFlex OS GUI	Alert Message in REST	Alert Message in SNMP Trap	Alert Code (for ESRS)	Severity	Recommended Action
License expired	LICENSE_EXPIRED	System.License.License_Expired	SIO01.02.0000001	5 (Critical)	To resume operational mode, contact EMC Support for license renewal. If you have already renewed your license, install it.
The system's license will expire in n days	LICENSE_ABOUT_TO_EXPIRE	System.License.License_Is_About_To_Expire	SIO01.02.0000002	3 (Error) 2 (Warning) according to	Contact EMC Support for license renewal. If you have

Table 33 VxFlex OS Alerts in SNMP, VxFlex OS GUI, REST, and ESRS (continued)

Alert Message in VxFlex OS GUI	Alert Message in REST	Alert Message in SNMP Trap	Alert Code (for ESRS)	Severity	Recommended Action
				time left and limits	already renewed your license, install it.
VxFlex OS is using a trial license	TRIAL_LICENSE_USED	System.License.Trial_License_Used	SIO01.02.0000003	2 (Warning)	Purchase a license and install it.
Oscillating failures reported	OBJECT_HAS_OSCILLATING_FAILURES	System.Oscillating_Failures.Object_has_oscillating_failures	SIO01.03.0000001	2 (Warning)	Check oscillating failures of the component and take action accordingly. If the oscillating failure does not indicate a problem, change the settings of the oscillating failure window to supress this alert
There are oscillating network failures	OBJECT_HAS_OSCILLATING_NETWORK_FAILURES	System.Oscillating_Failures.OBJECT_HAS_OSCILLATING_NETWORK_FAILURES	SIO01.03.0000002	2 (Warning)	Check the oscillating failure report, that can be accessed from one of the management interfaces. Check whether there is a problem with network links, fix, and restart the counters.
Invalid MDM monitoring credentials are configured in the AMS	AMS_CONFIGURATION_INVALID_MDM_CREDENTIALS	System.Credentials.Ams_Configuration_Invalid_Mdm_Credentials	SIO01.04.0000003	5 (Critical)	Check the MDM credentials configured in the AMS, and modify them if necessary
The AMS cannot connect to the MDM cluster, but the cluster itself is operational	AMS_LOST_CONNECTION_TO_MDM	System.MDM_Cluster.Ams_Lost_Connection_To_Mdm	SIO01.06.0000001	5 (Critical)	Switch MDM ownership using one of the management interfaces. This change may pass cluster control to a master MDM to which the AMS can establish a connection.
The MDM is not operating in clustered mode	MDM_NOT_CLUSTERED	MDM.MDM_Cluster.MDM_Not_Clustered	SIO02.01.0000001	5 (Critical)	MDM cluster was manually set to SINGLE mode. Confirm that this is an expected operation. Working in SINGLE mode is not recommended. Prepare the cluster modules (if needed), and return to CLUSTER mode.
MDM fails over frequently	MDM_FAILS_OVER_FREQUENTLY	MDM.MDM_Cluster.MDM_Fail	SIO02.01.0000003	5 (Critical) 3 (Error) 2 (Warning)	The MDMs frequently swap ownership. No action required.

Table 33 VxFlex OS Alerts in SNMP, VxFlex OS GUI, REST, and ESRS (continued)

Alert Message in VxFlex OS GUI	Alert Message in REST	Alert Message in SNMP Trap	Alert Code (for ESRS)	Severity	Recommended Action
		s_Over_Frequently		according to disconnect count and hardcoded values (2/3/10)	
Forward rebuild cannot proceed	FWD_REBUILD_STUCK	MDM.MDM_Cluster.FWD_REBUILD_STUCK	SIO02.01.0000004	2 (Warning)-5 (Critical)	Check the system for lack of spare capacity and/or failed capacity, and either fix the problem or add capacity if necessary.
Backward rebuild cannot proceed	BKWD_REBUILD_STUCK	MDM.MDM_Cluster.BKWD_REBUILD_STUCK	SIO02.01.0000005	2 (Warning)-5 (Critical)	Check the system for lack of spare capacity and/or failed capacity, and either fix the problem or add capacity if necessary.
Rebalance cannot proceed	REBALANCE_STUCK	MDM.MDM_Cluster.REBALANCE_STUCK	SIO02.01.0000006	5 (Critical) 3 (Error) 2 (Warning)	Add a physical disk; if this is not possible, reduce the spare policy while maintaining enough spare to sustain a rebuild, if necessary.
The MDM cluster is degraded, and data is not protected	CLUSTER_DEGRADED	MDM.MDM_Cluster.CLUSTER_DEGRADED	SIO02.01.0000007	3 (Error)-5 (Critical)	Check that all MDM cluster nodes are functioning correctly, and fix and replace faulty nodes, if necessary, in order to return to full protection.
Cannot connect to the MDM cluster, but the cluster itself is operational	MDM_CONNECTION_LOST	MDM.MDM_Cluster.MDM_CONNECTION_LOST	SIO02.01.0000008		Check the connection to the MDM
The MDM is not operating in Clustered mode	MDM_NOT_CLUSTERED_VOLUMES_EXIST	MDM.MDM_Cluster.MDM_Not_Clustered_Volume_Exist	SIO02.01.0000009	5 (Critical)	The MDM cluster was manually set to SINGLE mode. Working in SINGLE mode is not recommended. Single mode means that there is only one copy of the MDM repository. If you lose this copy, all System configurations and all the data on all the existing volumes will be lost. Please verify that this is an expected operation. Prepare the cluster

Table 33 VxFlex OS Alerts in SNMP, VxFlex OS GUI, REST, and ESRS (continued)

Alert Message in VxFlex OS GUI	Alert Message in REST	Alert Message in SNMP Trap	Alert Code (for ESRS)	Severity	Recommended Action
					modules (if needed), and return to CLUSTERED mode as soon as possible.
Inactive Protection Domain	PD_INACTIVE	MDM.Protection_Domain.Protection_Inactive	SIO02.02.0000001	2 (Warning)	Protection Domain was inactivated by a user command. Confirm that this is an expected operation. This is usually done for maintenance. When maintenance is complete, reactivate the Protection Domain.
Storage Pool has failed capacity	STORAGE_POOL_HAS_FAILED_CAPACITY	MDM.Storage_Pool.Storage_Pool_has_Failed_Capacity	SIO02.04.0000001	5 (Critical)	For the given Storage Pool, for some blocks, both primary and secondary copies are inaccessible. Check and fix the state of all devices in the Storage Pool and all the server's holding devices in the Storage Pool.
Storage Pool has degraded capacity	STORAGE_POOL_HAS_DEGRADED_CAPACITY	MDM.Storage_Pool.Storage_Pool_has_Degraded_Capacity	SIO02.04.0000002	3 (Error)	For the given Storage Pool, for some blocks, one of the two copies is inaccessble. Check if a server is offline or if there is another server hardware-related issue. Check if a storage device is down.
Capacity utilization above critical threshold	CAPACITY_UTILIZATION_ABOVE_CRITICAL_THRESHOLD	MDM.Storage_Pool.Capacity_Utilization_Above_Critical_Threshold	SIO02.04.0000003	5 (Critical)	Due to thinly provisioned volumes or snapshot usage, the capacity utilization of the Storage Pool is reaching a critical threshold. Remove snapshots, if possible, or add physical storage.
Capacity utilization above high threshold	CAPACITY_UTILIZATION_ABOVE_HIGH_THRESHOLD	MDM.Storage_Pool.Capacity_Utilization_Above_High_Threshold	SIO02.04.0000004	2 (Warning) 3 (Error)	Due to thinly provisioned volumes or snapshot usage, the capacity utilization of the Storage Pool is reaching a high threshold. Remove snapshots, if possible, or add physical storage.

Table 33 VxFlex OS Alerts in SNMP, VxFlex OS GUI, REST, and ESRS (continued)

Alert Message in VxFlex OS GUI	Alert Message in REST	Alert Message in SNMP Trap	Alert Code (for ESRS)	Severity	Recommended Action
Failure recovery capacity is below the threshold	FAILURE_RECOVERY_CAPACITY_BELOW_THRESHOLD	MDM.Storage_Pool.Failure_Recovery_Capacity_Below_Threshold	SIO02.04.0000005	3 (Error)	The capacity available for recovery in a degraded storage event is lower than the predefined threshold. Replace failed hardware or add more physical storage.
Configured spare capacity is smaller than largest fault unit	CONFIGURED_SPARE_CAPACITY_SMALLER_THAN_LARGEST_FAULT_UNIT	MDM.Storage_Pool.Configured_Spare_Capacity_Smaller_Than_Largest_Fault_Unit	SIO02.04.0000008	2 (Warning)	Increase the "spare percentage" configured in the system for the Storage Pool, so that the capacity reserved for failure recovery is larger than the largest fault unit in the Storage Pool.
The Storage Pool relies too heavily (over 50%) on capacity from a single SDS or Fault Set. Balance capacity over other SDSs or Fault Sets.	STORAGE_POOL_UNBALANCED	MDM.Storage_Pool.STORAGE_POOL_UNBALANCED	SIO02.04.0000009	3 (Error)	Move some physical disks from the large SDS to the others, or add disks to the smaller SDS in order to approximate the capacity of the large SDS as much as possible.
Storage Pool does not meet the minimum requirement of 3 fault units	NOT_ENOUGH_FAULT_UNITS_IN_SP	MDM.Storage_Pool.Not_Enough_Fault_Units	SIO02.04.0000010	3 (Error)	Add more SDSs to the Storage Pool to meet the minimum requirement of 3 hosts.
In the Storage Pool, some devices are not being cached (using DAS Cache), even though the Storage Pool is marked as cached	CACHECADE_NOT_USED_IN_STORAGE_POOL_THOUGH_IT_SHOULD	MDM.Storage_Pool.Das_Cache_Not_Used_In_Storage_Pool_Though_It_Should	SIO02.04.0000013	2 (Warning)	Check the caching configuration of those devices. Caching needs to be enabled at both device and Storage Pool levels.
In the Storage Pool, some devices are being cached (using DAS Cache) even though the Storage Pool is not marked as cached	CACHECADE_USED_IN_STORAGE_POOL_THOUGH_IT_SHOULDNT	MDM.Storage_Pool.Das_Cache_Used_In_Storage_Pool_Though_It_Shouldnt	SIO02.04.0000014	2 (Warning)	Stop DAS Cache on the Storage Pool.
There are cluster certificates pending approval. For more information, open	UNTRUSTED_CERTIFICATE	MDM.CERTIFICATE.UNTRUSTED_CERTIFICATE	SIO02.05.0000001	3 (Error)	Pending for approval certificates can be viewed and approved via System Settings > Connection.

Table 33 VxFlex OS Alerts in SNMP, VxFlex OS GUI, REST, and ESRS (continued)

Alert Message in VxFlex OS GUI	Alert Message in REST	Alert Message in SNMP Trap	Alert Code (for ESRS)	Severity	Recommended Action
System Settings > Connection.					
Master MDM Certificate is about to expire	CERTIFICATE_ABOUT_TO_EXPIRE	MDM.CERTIFICATE.CERTIFICATE_ABOUT_TO_EXPIRE	SIO02.05.0000002	5 (Critical)	Install a valid SSL certificate on the MDM before the old one expires
Master MDM Certificate has expired	MDM_CERTIFICATE_EXPIRED	MDM.CERTIFICATE.MDM_CERTIFICATE_EXPIRED	SIO02.05.0000003	5 (Critical)	Install a valid SSL certificate on the host
Secure connection disabled on MDM	MDM_SECURE_CONNECTION_DISABLED	MDM.CERTIFICATE.MDM_SECURE_CONNECTION_DISABLED	SIO02.05.0000004	5 (Critical)	Enable secure connections on the MDM in order to protect your login information
The self-signed certificate presented by the Master MDM is not trusted	MDM_SELF_SIGNED_CERTIFICATE_NOT_TRUSTED	MDM.CERTIFICATE.MDM_SELF_SIGNED_CERTIFICATE_NOT_TRUSTED	SIO02.05.0000005	5 (Critical)	Check the certificate, and trust it if you see fit to do so
MDM does not support secure connections	MDM_SECURE_CONNECTION_NOT_SUPPORTED	MDM.CERTIFICATE.MDM_SECURE_CONNECTION_NOT_SUPPORTED	SIO02.05.0000006	5 (Critical)	Check MDM cluster nodes
The validity period of the certificate presented by the Master MDM starts in the future	MDM_CERTIFICATE_NOT_YET_VALID	MDM.CERTIFICATE.MDM_CERTIFICATE_NOT_YET_VALID	SIO02.05.0000007	5 (Critical)	The time and date on the computer where the certificate was created is not consistent with the time and date set in the VxFlex OS system. Replace the certificate or fix the system time.
The Certificate Authority that signed the Master MDM's certificate is not trusted	MDM_CA_SIGNED_CERTIFICATE_CA_NOT_TRUSTED	MDM.CERTIFICATE.MDM_CA_SIGNED_CERTIFICATE_CA_NOT_TRUSTED	SIO02.05.0000008	5 (Critical)	Trust the CA certificate if you see fit
SDS is disconnected	SDS_DISCONNECTED	SDS.SDS.SDS_Disconnected	SIO03.01.0000001	3 (Error)	The SDS service may be down or unreachable over the network. Verify that the SDS service is up and running and that the network is properly connected.

Table 33 VxFlex OS Alerts in SNMP, VxFlex OS GUI, REST, and ESRS (continued)

Alert Message in VxFlex OS GUI	Alert Message in REST	Alert Message in SNMP Trap	Alert Code (for ESRS)	Severity	Recommended Action
SDS disconnects frequently	SDS_DISCONNECTS_FREQUENTLY	SDS.SDS.SDS_Disconnected_Frequently	SIO03.01.0000002	3 (Error) 2 (Warning) according to disconnect count and hard-coded	The SDS connection is fluctuating due to an unstable network connection. Check the SDS data network connection for Packet Drops, and try to disconnect one of the ports to see if the SDS disconnection issue is resolved by using only one port. If this does not resolve the issue, switch to the other port. If there is still an issue, it may be due to a Faulty NIC, Faulty Switch ports, or a faulty switch. If there is no issue with another switch, the issue was switch-related. Otherwise, the issue may be due to a faulty NIC, which requires NIC replacement.
Memory allocation for RAM Read Cache failed on SDS	SDS_RMCACHE_MEMORY_ALLOCATION_FAILED	SDS.SDS.SDS_Rmcache_Memory_allocation_Failed	SIO03.01.0000003	2 (Warning)	The system failed to allocate memory to the SDS RAM Read-Cache. For 32 GB RAM or less, up to 50% of the memory can be allocated for caching. From 32 GB or more, up to 75% of the memory can be allocated for caching. Reduce the configured RAM Read-Cache memory to match the allocation conditions.
DRL mode: Hardened	DRL_MODE_NON_VOLATILE	SDS.SDS.DRL_MODE_NON_VOLATILE	SIO03.01.0000004	1 (Info)	DRL Mode is configured to "Hardened" instead of "Volatile". Both modes are configurable.
RFcache card I/O error	RFCACHE_CARD_IO_ERROR	SDS.SDS.RFCACHE_CARD_IO_ERROR	SIO03.01.0000005	2 (Warning)	Disable caching on the device and check the health of the device, because it may be faulty. If necessary, replace the device.

Table 33 VxFlex OS Alerts in SNMP, VxFlex OS GUI, REST, and ESRS (continued)

Alert Message in VxFlex OS GUI	Alert Message in REST	Alert Message in SNMP Trap	Alert Code (for ESRS)	Severity	Recommended Action
RFcache skipped due to heavy load	RFCACHE_CACHE_SKIP_DUE_TO_HEAVY_LOAD	SDS.SDS.RFCACHE_CACHE_SKIP_DUE_TO_HEAVY_LOAD	SIO03.01.0000006	2 (Warning)	Read Flash Cache is working under a heavy load, and therefore has skipped some IOs. This is a temporary error which should resolve itself. If it persists, try to balance the Storage Pool contents across more SDSs, or add more cache cards.
RFcache IO stack error	RFCACHE_IO_STUCK_ERROR	SDS.SDS.RFCACHE_IO_STUCK_ERROR	SIO03.01.0000007	2 (Warning)	IO has become stuck on the cache device. Disable caching on the device and check the health of the device, because it may be faulty. If necessary, replace the device.
RFcache resources are low	RFCACHE_LOW_RESOURCES	SDS.SDS.RFCACHE_LOW_RESOURCES	SIO03.01.0000008	2 (Warning)	There is not enough RAM available on the server for Read Flash Cache optimal operation. Increase the amount of available RAM.
RFcache driver path is invalid	RFCACHE_INVALID_DRIVER_PATH	SDS.SDS.RFCACHE_INVALID_DRIVER_PATH	SIO03.01.0000009	2 (Warning)	The Read Flash cache driver (xcache) is either not installed, or was installed in the wrong location. Install the driver, and contact Customer Support if the problem persists.
RFcache source configuration is not consistent	RFCACHE_INCONSISTENT_SOURCE_CONFIGURATION	SDS.SDS.RFCACHE_INCONSISTENT_SOURCE_CONFIGURATION	SIO03.01.0000010	2 (Warning)	Check RFcache state of all disks in the pool and adjust them so that all disks have the same caching state.
RFcache source configuration is not consistent	RFCACHE_INCONSISTENT_CACHE_CONFIGURATION	SDS.SDS.RFCACHE_INCONSISTENT_CACHE_CONFIGURATION	SIO03.01.0000011	2 (Warning)	Query the system to determine what is not consistent in the configurations of the Read Flash cache driver and the SDS where the cache device is located.
RFcache device does not exist	RFCACHE_DEVICE_DOES_NOT_EXIST	SDS.SDS.RFCACHE_DEVICE_DOES_NOT_EXIST	SIO03.01.0000012	2 (Warning)	You tried to add a cache device that does not exist. Check and fix Read Flash Cache configuration.

Table 33 VxFlex OS Alerts in SNMP, VxFlex OS GUI, REST, and ESRS (continued)

Alert Message in VxFlex OS GUI	Alert Message in REST	Alert Message in SNMP Trap	Alert Code (for ESRS)	Severity	Recommended Action
RFcache API mismatch	RFCACHE_API_ERROR_MISMACH	SDS.SDS.RFCACHE_API_ERROR_MISMACH	SIO03.01.0000013	2 (Warning)	The Read Flash Cache (xcache) driver version and SDS version do not match. Try to upgrade them to the same version. If the problem persists, contact Customer Support.
SDS is in Maintenance Mode	SDS_IN_MAINTENANCE	SDS.SDS.SDS_IN_MAINTENANCE	SIO03.01.0000014	2 (Warning)	The SDS is currently in Maintenance Mode. Exit Maintenance Mode once maintenance is complete. If an NDU is in progress, ignore this warning.
Device failed	DEVICE_FAILED	SDS.Device.Device_Failed	SIO03.02.0000001	3 (Error)	The SDS device could not be opened, read from or written to. Validate the device state. Check the cause of the error, and determine if it is a human error or a system malfunction. Check hardware if needed.
Device test is done and device is pending activation	DEVICE_PENDING_ACTIVATION	SDS.Device.Device_Pending_Activation	SIO03.02.0000002	2 (Warning)	The SDS device has been added and tested. Activate the SDS device.
Device has fixed read errors	FIXED_READ_ERROR_COUNT_ABOVE_WARNING_THRESHOLD	SDS.Device.FIXED_READ_ERROR_COUNT_ABOVE_WARNING_THRESHOLD	SIO03.02.0000003	3 (Error) if counter > 0	Read from the SDS device failed. Data was corrected from the other copy. No action is required, but note that the device might be faulty.
Device has fixed read errors	FIXED_READ_ERROR_COUNT_ABOVE_CRITICAL_THRESHOLD	SDS.Device.FIXED_READ_ERROR_COUNT_ABOVE_CRITICAL_THRESHOLD	SIO03.02.0000004	5 (Critical) if counter >= 5	SDS device read failed more than 5 times. Replace the physical device.
Device failed: All IO to the device will be stopped, and data will be relocated to another device.	DEVICE_ERROR_ERROR	SDS.Device.Device_ErrorError	SIO03.02.0000005	5 (Critical)	Check the device, and if necessary, replace it
Device failed: All IO to the device will be stopped, and data	DEVICE_ERROR_WARNING	SDS.Device.Device_ErrorWarning	SIO03.02.0000006	3 (Error)	Check the device, and if necessary, replace it

Table 33 VxFlex OS Alerts in SNMP, VxFlex OS GUI, REST, and ESRS (continued)

Alert Message in VxFlex OS GUI	Alert Message in REST	Alert Message in SNMP Trap	Alert Code (for ESRS)	Severity	Recommended Action
will be relocated to another device.					
Device malfunction has been detected.	DEVICE_ERROR_NOTICE	SDS.Device.Device_ErrorNotice	SIO03.02.0000007	2 (Warning)	Check the device, and if necessary, replace it
Minor failures have been detected in device performance.	DEVICE_ERROR_INFO	SDS.Device.Device_ErrorInfo	SIO03.02.0000008	2 (Warning)	Check the device, and if necessary, replace it
The SDC is either down or unreachable over the network	SDC_DISCONNECTED	SDC.SDC.SDC_DISCONNECTED	SIO04.01.0000001	3 (Error)	Verify that the SDC service is up and running and that the network is properly configured and connected.
No more SDCs can be defined on this system; the maximum has been reached	SDC_MAX_COUNT	SDC.SDC.SDC_Max_Count	SIO04.01.0000002	3 (Error)	The maximum number of SDCs in the system has been reached (1024^2).
SSD lifespan is over endurance threshold	PHYSICAL_DRIVE_ENDURANCE_USED_ABOVE_THRESHOLD	Storage_Controller.Physical_Drive.Endurance_Used_Above_Threshold	SIO06.01.0000004	5 (Critical)	Replace the SSD device.
The Physical Disk temperature has exceeded the configured threshold	PHYSICAL_DRIVE_INVALID_TEMPERATURE	Storage_Controller.Physical_Drive.INVALID_TEMPERATURE	SIO06.01.0000005	5 (Critical) 3 (Error)	Ensure that the server and its environment are properly cooled. Ensure that the internal chassis fans are working, and that there is adequate airflow.
The Physical Drive's cache is not disabled	PHYSICAL_DRIVE_INVALID_PD_CACHE_POLICY	Storage_Controller.Physical_Drive.INVALID_PD_CACHE_POLICY	SIO06.01.0000006	3 (Error)	Disable the Physical Drive physical cache. Note: This may cause DI issues in power cycle scenarios.
Device was requested to use as DAS Cache but it is not	DEVICE_SHOULD_USE_AS_DAS_CACHE_BUT_IT_IS_NOT	Storage_Controller.Physical_Drive.Device_Should_Use_As_Das_Cache_But_It_Is_Not	SIO06.01.0000007	3 (Error)	
The Logical Disk read policy is not set to "Read-Ahead"	LOGICAL_DISK_INVALID_READ_AHEAD_POLICY	Storage_Controller.Logical_Disk.INVALID_READ_AHEAD_POLICY	SIO06.02.0000001	3 (Error)	Access the storage controller and change the Logical Disk's read policy to "Read Ahead".

Table 33 VxFlex OS Alerts in SNMP, VxFlex OS GUI, REST, and ESRS (continued)

Alert Message in VxFlex OS GUI	Alert Message in REST	Alert Message in SNMP Trap	Alert Code (for ESRS)	Severity	Recommended Action
The Logical Disk IO policy is not Configured properly (Write-Back)	LOGICAL_DISK_INVALID_WRITE_BACK_POLICY	Storage_Controller.Logical_Disk.INVALID_WRITE_BACK_POLICY	SIO06.02.0000002	3 (Error)	Check the Raid controller battery status, a fault in the battery will impact the Write-Back functionality.
Logical Disk access mode is not set to "Read-Write"	LOGICAL_DISK_INVALID_ACCESS_MODE	Storage_Controller.Logical_Disk.INVALID_ACCESS_MODE	SIO06.02.0000003	5 (Critical)	Access the storage controller and change the logical disk's access mode to "Read-Write".
The Logical Disk RAID level is not set to RAID0	LOGICAL_DISK_INVALID_RAID_LEVEL	Storage_Controller.Logical_Disk.INVALID_RAID_LEVEL	SIO06.02.0000004	3 (Error)	The RAID type is set in the storage controller is incorrect. Access the storage controller, verify that the Logical Disk does not contain any data that is not backed-up, destroy the Logical Disk, and re-create the logical disk as RAID0.
Logical Disk caching policy is not set to "DirectIO"	LOGICAL_DISK_INVALID_CACHE_POLICY	Storage_Controller.Logical_Disk.INVALID_CACHE_POLICY	SIO06.02.0000005	3 (Error)	Access the storage controller, and change the RAID0 caching policy for the Logical Disk to "DirectIO".
Logical disk is no longer being cached	LOGICAL_DISK_NO_LONGER_CACHED	Storage_Controller.Logical_Disk.LOGICAL_DISK_NO_LONGER_CACHED	SIO06.02.0000006	3 (Error)	Disable the Physical Drive physical cache. Note: This may cause DI issues in power cycle scenarios.
The state of the backup battery in the storage controller is not optimal	BATTERY_INVALID_STATE	Storage_Controller.Battery.Invalid_State	SIO06.03.0000001	5 (Critical) 3 (Error)	The backup battery is not fully charged, but it will recharge itself while the storage controller is powered on.
The backup battery in the storage controller needs to be replaced	BATTERY_REPLACEMENT_REQUIRED	Storage_Controller.Battery.REPLACEMENT_REQUIRED	SIO06.03.0000002	5 (Critical)	The backup battery in the storage controller may be nearing the end of its working life. Replace the battery.
The storage controller battery temperature is not within the configured threshold	BATTERY_INVALID_TEMPERATURE	Storage_Controller.Battery.INVALID_TEMPERATURE	SIO06.03.0000003	3 (Error)	Check the temperature in the server and at the data center. Check if a fan alert is raised in the node. Check the battery and the RAID controller and replace faulty items.

Table 33 VxFlex OS Alerts in SNMP, VxFlex OS GUI, REST, and ESRS (continued)

Alert Message in VxFlex OS GUI	Alert Message in REST	Alert Message in SNMP Trap	Alert Code (for ESRS)	Severity	Recommended Action
There is no backup battery installed in the storage controller	BATTERY_NOT_PRESENT	Storage_Controller.Battery.NOT_PRESENT	SIO06.03.0000004	5 (Critical)	Install a backup battery in the storage controller.
The storage controller contains an incompatible battery pack	BATTERY_INVALID_PACK_ENERGY	Storage_Controller.Battery.Invalid_Pack_Energy	SIO06.03.0000005	3 (Error)	Replace the storage controller battery with one that is compatible with your controller's model.
RAID controller's battery has invalid voltage state	BATTERY_INVALID_VOLTAGE	Storage_Controller.Battery.Invalid_Voltage	SIO06.03.0000006	3 (Error)	Check the RAID controller's battery. It may need to be replaced.
Hypervisor boot drive state is not OK	BOOT_DRIVE_INVALID_STATE	Node.Boot_Drive.Invalid_State	SIO07.01.0000001	3 (Error)	The drive on which the ESXi hypervisor is installed has reported an error. Replace the drive and reinstall the ESXi on the new drive.
The storage controller is not operating optimally	STORAGE_CONTROLLER_INVALID_STATE	Node.Storage_Controller.Invalid_State	SIO07.02.0000001	5 (Critical)	The storage controller may be faulty, and should be replaced.
The storage controller temperature has exceeded the configured threshold	STORAGE_CONTROLLER_INVALID_TEMPERATURE	Node.Storage_Controller.Invalid_Temperature	SIO07.02.0000002	5 (Critical) 3 (Error)	Ensure that the server and its environment are properly cooled and that all the chassis fans are functional.
CacheCade license is not installed	STORAGE_CONTROLLER_CACHECADE_NOT_LICENSED	Node.Storage_Controller.Cachecade_Not_Licensed	SIO07.02.0000003	3 (Error)	Install a CacheCade license on the storage controller.
Some of the disk slots in the storage controller are empty	STORAGE_CONTROLLER_NOT_ALL_SLOTS_FULL	Node.Storage_Controller.Not_All_Slots_Full	SIO07.02.0000004	3 (Error)	Add more disks if needed. If a disk was removed during a FRU (Field Replacement Unit) operation, insert the new disk into the chassis.
Unable to query RAID controller due to insufficient permissions	NODE_FAILED_TO_QUERY_RAID_CONTROLLER_INVALID_PERMISSIONS	Node.Storage_Controller.NODE_FAILED_TO_QUERY_RAID_CONTROLLER_INVALID_PERMISSIONS	SIO07.02.0000005	5 (Critical)	Verify that the user name and password of the Storage VM match the user name and password configured in the AMS (System Settings menu > System Settings > Security option). Resolve the differences and try again.

Table 33 VxFlex OS Alerts in SNMP, VxFlex OS GUI, REST, and ESRS (continued)

Alert Message in VxFlex OS GUI	Alert Message in REST	Alert Message in SNMP Trap	Alert Code (for ESRS)	Severity	Recommended Action
The physical CPU socket is not enabled	SOCKET_DISABLED	Node.Cpu_Socket.Disabled	SIO07.03.0000001	3 (Error)	Enable the CPU socket in the server BIOS.
Not all CPU cores are enabled	SOCKET_NOT_ALL_CORES_ENABLED	Node.Cpu_Socket.NOT_ALL_CORES_ENABLED	SIO07.03.0000002	3 (Error)	Change the physical CPU core configuration in the server BIOS.
CPU is not operating at full speed	SOCKET_SPEED_IS_NOT_MAX_SPEED	Node.Cpu_Socket.SPEED_IS_NOT_MAX_SPEED	SIO07.03.0000003	3 (Error)	Set the CPU clock speed in the server's BIOS to its maximum speed.
CPU temperature is not within threshold	CPU_SOCKET_INVALID_TEMPERATURE	Node.Cpu_Socket.INVALID_TEMPERATURE	SIO07.03.0000004	5 (Critical) 3 (Error)	The server's CPU temperature has exceeded the configured threshold. Make sure that the server is properly cooled and that the CPU and internal chassis fans are active.
The CPU Voltage Regulator's temperature is higher than the configured threshold	CPU_SOCKET_INVALID_VR_TEMPERATURE	Node.Cpu_Socket.INVALID_VR_TEMPERATURE	SIO07.03.0000005	5 (Critical) 3 (Error)	Make sure that the server is properly cooled and that the CPU and internal chassis fans are active.
The CPU Voltage Regulator's voltage is not within the configured threshold	CPU_SOCKET_INVALID_VR_VOLTAGE	Node.Cpu_Socket.INVALID_VR_VOLTAGE	SIO07.03.0000006	5 (Critical) 3 (Error)	1. Verify that the power supply is functioning correctly. 2. Try to replace a port in the Power Distribution Unit, or supply an external power source to check it. 3. Replace the power cable. 4. Replace the Power Supply Unit module.
CPU has invalid voltage state state	CPU_SOCKET_INVALID_VOLTAGE	Node.Cpu_Socket.Invalid_Voltage	SIO07.03.0000007	5 (Critical)	Check the chassis power supply
RAM temperature is not within threshold	RAM_INVALID_TEMPERATURE	Node.Ram.Invalid_Temperature	SIO07.04.0000001	5 (Critical) 3 (Error)	The temperature of one or more Server RAM modules exceeds the configured threshold. Ensure that the server and its environment are properly cooled and that all the chassis fans are functional.

Table 33 VxFlex OS Alerts in SNMP, VxFlex OS GUI, REST, and ESRS (continued)

Alert Message in VxFlex OS GUI	Alert Message in REST	Alert Message in SNMP Trap	Alert Code (for ESRS)	Severity	Recommended Action
RAM voltage regulator's temperature is not within threshold	RAM_INVALID_VR_TEMPERATURE	Node.Ram.INVALID_VR_TEMPERATURE	SIO07.04.0000002	5 (Critical) 3 (Error)	Ensure that the server and its environment are properly cooled and that all the chassis fans are functional.
RAM voltage regulator's voltage is not within threshold	RAM_INVALID_VR_VOLTAGE	Node.Ram.INVALID_VR_VOLTAGE	SIO07.04.0000003	5 (Critical) 3 (Error)	Ensure that the server and its environment are properly cooled and that all the chassis fans are functional. If the system still issues an alert, the DIMM may be faulty, and may have to be replaced.
The node temperature has exceeded the configured threshold	NODE_INVALID_TEMPERATURE	Node.Node.INVALID_TEMPERATURE	SIO07.05.0000001	5 (Critical) 3 (Error)	Ensure that the server and its environment are properly cooled and that all the chassis fans are functional.
The voltage values on the node are not within the configured threshold	NODE_INVALID_VOLTAGE	Node.Node.INVALID_VOLTAGE	SIO07.05.0000002	5 (Critical) 3 (Error)	1. Verify that the power supply is functioning correctly, and then try to replace a port in the Power Distribution Unit or check with external power. 2. Replace the power cable. 3. Replace the Power Supply Module.
Unable to connect to monitoring agent (BMC)	NODE_FAILED_TO_CONNECT_TO_BMC	Node.Node.FAILED_TO_CONNECT_TO_BMC	SIO07.05.0000004	5 (Critical)	The BMC IP address or User credentials do not respond to AMS queries. Verify network connectivity from the AMS to the BMC IP address, and check that the BMC admin user and password have not been tampered with. Contact customer support for assistance if the problem persists.
SDC is not installed on this node	NODE_WITH_NO_SDC	Node.Node.NODE_WITH_NO_SDC	SIO07.05.0000005	3 (Error)	Consider installing an SDC on this node, so that it can use VxFlex OS volumes
Unable to connect to monitoring agent (SVM)	NODE_FAILED_TO_CONNECT_TO_SVM	Node.Node.FAILED_TO_CONNECT_TO_SVM	SIO07.05.0000006	5 (Critical)	The SVM IP address or User credentials do not respond to AMS queries. Check the SVM state in

Table 33 VxFlex OS Alerts in SNMP, VxFlex OS GUI, REST, and ESRS (continued)

Alert Message in VxFlex OS GUI	Alert Message in REST	Alert Message in SNMP Trap	Alert Code (for ESRS)	Severity	Recommended Action
					vCenter and try to open the SSH or the console to it. Then check the message log for errors.
Unable to connect to monitoring agent (ESX)	NODE_FAILED_TO_CONNECT_TO_ESX	Node.Node.FAILED_TO_CONNECT_TO_ESX	SIO07.05.0000007	5 (Critical)	The ESX IP address or User credentials do not respond to AMS queries. Check the ESX state in vCenter and try to open the SSH or the console (via BMC http) to it. Then check the messages log for errors.
Unable to connect to monitoring agent (host)	NODE_FAILED_TO_CONNECT_TO_HOST	Node.Node.FAILED_TO_CONNECT_TO_HOST	SIO07.05.0000007	5 (Critical)	The ESX IP address or User credentials do not respond to AMS queries. Check the ESX state in vCenter and try to open the SSH or the console (via BMC http) to it. Then check the messages log for errors.
Unable to connect to monitoring agent (VCenter)	NODE_FAILED_TO_CONNECT_TO_VCENTER	Node.Node.FAILED_TO_CONNECT_TOVCENTER	SIO07.05.0000008	5 (Critical)	Ensure that the vCenter configuration is correct, and ensure that the node's Mgmt port is routable to the vCenter IP address
Node serial number has changed	NODE_SERIAL_NUMBER_CHANGE	Node.Node.SERIAL_NUMBER_CHANGE	SIO07.05.0000009	3 (Error)	The S/N of the Motherboard does not match the IP address of the ESX. Either the SATADOM was moved to a new server or the motherboard was replaced. It is also possible that a new server that has the same IP addresses and the same user and password, but was not installed using the AMS flow, exists in the network. EMC does not support such field replacement cases. Such cases should be solved by replacing nodes using EMC-recommended procedures. For more information, contact Customer Support.

Table 33 VxFlex OS Alerts in SNMP, VxFlex OS GUI, REST, and ESRS (continued)

Alert Message in VxFlex OS GUI	Alert Message in REST	Alert Message in SNMP Trap	Alert Code (for ESRS)	Severity	Recommended Action
RFcache pool state is "Not Started"	XTREMCACHE_INVALID_STATE	Node.Node.XTREMCACHE_INVALID_STATE	SIO07.05.0000010	3 (Error)	Check the Read Flash Cache pool state. The cache device might be misconfigured. In this case, remove the cache device from the cache pool and add it back again.
Host certificate is about to expire	HOST_CERTIFICATE_ABOUT_TO_EXPIRE	Node.Node.HOST_CERTIFICATE_ABOUT_TO_EXPIRE	SIO07.05.0000014	3 (Error)	Renew the AMS certificate and then run "renew certificate" process
CMOS battery state is invalid	NODE_INVALID_CMOS_BATTERY	Node.Node.Invalid_Cmos_Battery	SIO07.05.0000015	3 (Error)	Check the node's CMOS battery. It may need to be replaced.
The node is not being managed by the current AMS.	NODE_NOT_MANAGED_BY_AMS	Node.Node.Node_Not_Managed_By_Ams	SIO07.05.0000016	5 (Critical)	Verify that the node is not being managed by a different AMS.
Failed to connect to the SVM due to authentication failure	NODE_FAILED_TO_CONNECT_TO_SVM_AUTH_FAILED	Node.Node.Node_Failed_To_Connect_To_Svm_Auth_Failed	SIO07.05.0000017	5 (Critical)	Change the SVM's root password to be the one set by the AMS
Failed to connect to the host. The SSH key has been changed, and is not trusted by the AMS.	NODE_FAILED_TO_CONNECT_TO_HOST_SSH_KEY_MISMATCH	Node.Node.Node_Failed_To_Connect_To_Host_Ssh_Key_Mismatch	SIO07.05.0000018	5 (Critical)	This is a security issue. Remove the node, and add the node again to the AMS.
Failed to connect to the SVM. The SSH key has been changed, and is not trusted by the AMS.	NODE_FAILED_TO_CONNECT_TO_SVM_SSH_KEY_MISMATCH	Node.Node.Node_Failed_To_Connect_To_Svm_Ssh_Key_Mismatch	SIO07.05.0000019	5 (Critical)	This is a security issue. Remove the node, and add the node again to the AMS.
Failed to connect to the host because the certificate is not trusted	NODE_FAILED_TO_CONNECT_TO_ESX_CERTIFICATE_NOT_TRUSTED	Node.Node.Node_Failed_To_Connect_To_Esx_Certificate_Not_Trusted	SIO07.05.0000020	5 (Critical)	This is a security issue. Remove the node, and add the node again to the AMS.
Failed to connect to the host due to authentication failure	NODE_FAILED_TO_CONNECT_TO_HOST_AUTH_FAILED	Node.Node.Node_Failed_To_Connect_To_Host_Auth_Failed	SIO07.05.0000021	5 (Critical)	Change the host's root password to be the one set by the AMS
The MDM has a self-signed certificate that was not replaced by the	AMS_MDM_HAS_SELF_SIGNED_CERTIFICATE	Node.Node.Ams_Mdm_Has_Self_signed_certificate	SIO07.05.0000022	5 (Critical)	Disable AMS secure connection, and enable it again in order to automatically sign the MDM

Table 33 VxFlex OS Alerts in SNMP, VxFlex OS GUI, REST, and ESRS (continued)

Alert Message in VxFlex OS GUI	Alert Message in REST	Alert Message in SNMP Trap	Alert Code (for ESRS)	Severity	Recommended Action
AMS. Connection to the MDM is not authenticated by the AMS.					certificate. If the previous solution does not work, remove and add the the MDM again in secure mode.
CPU cache is not enabled	SOCKET_CACHE_DISABLED	Cpu_Socket.Socket_Cache.Disabled	SIO08.01.0000001	3 (Error)	Enable the CPU cache in the server BIOS.
CPU cache size is not optimal	SOCKET_CACHE_SIZE_NOT_MAX_SIZE	Cpu_Socket.Socket_Cache.SIZE_NOT_MAX_SIZE	SIO08.01.0000002	3 (Error)	Check the CPU cache in the server BIOS, and set it to maximum size.
Unable to connect to ESRS Gateway	ESRS_CONNECTIVITY_ERROR	Esrs.Esrs.CONNECTIVITY_ERROR	SIO10.01.0000001	5 (Critical)	Check the network's connectivity to the ESRS Gateway.
The system is not registered with an ESRS gateway	ESRS_NOT_REGISTERED	Esrs.Esrs.NOT_REGISTERED	SIO10.01.0000004	2 (Warning)	The system is not registered with an ESRS Gateway and will not send Alerts to EMC for monitoring. Contact your Support \ Sales representative to get an EMC ESRS support package.
ESRS number of messages received has been reached. No more messages will be sent in the next 8 hours	ESRS_REACHED_CAPACITY_LIMIT	Esrs.Esrs.REACHED_CAPACITY_LIMIT	SIO10.01.0000005	5 (Critical)	ESRS has a limit of receiving up to 200 alerts per 8 hours. The limit has been reached, so no ESRS messages will be sent in the following 8 hours.
SDC is disconnected from SDS	ONE_SDC_DISCONNECTED_FROM_ONE_SDS	System.SDC.One_Sdc_Disconnected_From_One_Sds	SIO01.07.0000001	5 (Critical)	Check the network links between the affected SDS and SDC
SDC is disconnected from the SDS IP address	ONE_SDC_DISCONNECTED_FROM_ONE_SDS_IP	System.SDC.One_Sdc_Disconnected_From_One_Sds_Ip	SIO01.07.0000002	2 (Warning)	Check the network links between the affected SDC and SDS IP addresses.
An SDC is disconnected from all SDSs	ONE_SDC_DISCONNECTED_FROM_ALL_SDS	System.SDC.One_Sdc_Disconnected_From_All_Sds	SIO01.07.0000003	5 (Critical)	Check the network links between the affected SDC and all SDSs.
All SDCs are disconnected from the SDS	ALL_SDC_DISCONNECTED_FROM_ONE_SDS	System.SDC.All_Sdc_Disconne	SIO01.07.0000004	5 (Critical)	Check the network links between all SDCs and the affected SDS.

Table 33 VxFlex OS Alerts in SNMP, VxFlex OS GUI, REST, and ESRS (continued)

Alert Message in VxFlex OS GUI	Alert Message in REST	Alert Message in SNMP Trap	Alert Code (for ESRS)	Severity	Recommended Action
		cted_From_One_Sds			
All SDCs are disconnected from the SDS IP address	ALL_SDC_DISCONNECTED_FROM_ONE_SDS_IP	System.SDC.All_Sdc_Disconnected_From_One_Sds_Ip	SIO01.07.0000005	2 (Warning)	Check the network links between all SDCs and the affected SDS IP address.
All SDCs are disconnected from all SDSs	ALL_SDC_DISCONNECTED_FROM_ALL_SDS	System.SDC.All_Sdc_Disconnected_From_All_Sds	SIO01.07.0000006	5 (Critical)	Check the network links between all SDCs and SDSs.
Disconnected network links between SDCs and SDSs.	SDC_MULTIPLE_DISCONNECTIONS_FROM_SDS	System.SDC.Sdc_Multiple_Disconnections_From_Sds	SIO01.07.0000007	5 (Critical)	Check the network links between all SDCs and SDSs.
This SDC has not been approved	SDC_NOT_APPROVED	System.SDC.Sdc_Not_Approved	SIO01.07.0000008	2 (Warning)	Add the IP address for this SDC to the list of approved IP addresses.
This SDC does not have an approved IP address	SDC_DOES_NOT_HAVE_APPROVED_IPS	System.SDC.Sdc_Does_Not_Have_Approved_Ips	SIO01.07.0000009	2 (Warning)	Update the list of approved IP addresses for this SDC.
SDC has an IP address that is not approved by the MDM	SDC_HAS_UNAPPROVED_IP	System.SDC.Sdc_Has_Unapproved_Ips	SIO01.07.0000010	3 (Error)	Add the IP address for this SDC to the list of approved IP addresses.
Some of the the approved SDC IP addresses in the MDM are not recognized by AMS	SDC_HAS_UNKNOWN_APPROVED_IP	System.SDC.Sdc_Has_Unknown_Approved_Ips	SIO01.07.0000011	2(Warning)	Right click the relevant SDC, select "Configure Approved IP Addresses", and add the second data link IP address

VxFlex Ready Node–specific Alerts in SNMP, VxFlex OS GUI, REST, and ESRS

This table summarizes additional alerts generated by systems with AMS support.

Table 34 VxFlex Ready Node–specific Alerts in SNMP, VxFlex OS GUI, REST, and ESRS

Alert Message in VxFlex OS GUI	Alert Message in REST	Alert Message in SNMP Trap	Alert Code (for ESRS)	Severity	Recommended Action
Errors have been detected on the Physical Disk	PHYSICAL_DRIVE_BAD_STATE	Storage_Controller.Physical_Drive.BAD_STATE	SIO06.01.0000001	5 (Critical)	Replace the Physical Disk.

Table 34 VxFlex Ready Node–specific Alerts in SNMP, VxFlex OS GUI, REST, and ESRS (continued)

Alert Message in VxFlex OS GUI	Alert Message in REST	Alert Message in SNMP Trap	Alert Code (for ESRS)	Severity	Recommended Action
Physical Disk state is erroneous (Hot Spare)	PHYSICAL_DRIVE_HOT_SPARE	Storage_Controller.Physical_Drive.HOT_SPARE	SIO06.01.0000002	3 (Error)	Errors have been detected on a spare Physical Disk. Replace the Physical Disk.
The Physical Disk is not being used by the system	PHYSICAL_DRIVE_NOT_IN_USE	Storage_Controller.Physical_Drive.NOT_IN_USE	SIO06.01.0000003	3 (Error)	Access the storage controller and configure the disk.
SSD lifespan is over endurance threshold	PHYSICAL_DRIVE_ENDURANCE_USED_ABOVE_THRESHOLD	Storage_Controller.Physical_Drive.Endurance_Used_Above_Threshold	SIO06.01.0000004	5 (Critical)	Replace the SSD device.
The Physical Disk temperature has exceeded the configured threshold	PHYSICAL_DRIVE_INVALID_TEMPERATURE	Storage_Controller.Physical_Drive.INVALID_TEMPERATURE	SIO06.01.0000005	5 (Critical) 3 (Error)	Ensure that the server and its environment are properly cooled. Ensure that the internal chassis fans are working, and that there is adequate airflow.
The Physical Drive's cache is not disabled	PHYSICAL_DRIVE_INVALID_PD_CACHE_POLICY	Storage_Controller.Physical_Drive.INVALID_PD_CACHE_POLICY	SIO06.01.0000006	3 (Error)	Disable the Physical Drive physical cache. Note: This may cause DI issues in power cycle scenarios.
Device was requested to use as DAS Cache but it is not	DEVICE_SHOULD_USE_AS_DAS_CACHE_BUT_IT_IS_NOT	Storage_Controller.Physical_Drive.Device_Should_Use_As_Das_Cache_But_It_Is_Not	SIO06.01.0000007	3 (Error)	
The Logical Disk read policy is not set to "Read-Ahead"	LOGICAL_DISK_INVALID_READ_AHEAD_POLICY	Storage_Controller.Logical_Disk.INVALID_READ_AHEAD_POLICY	SIO06.02.0000001	3 (Error)	Access the storage controller and change the Logical Disk's read policy to "Read Ahead".
The Logical Disk IO policy is not Configured properly (Write-Back)	LOGICAL_DISK_INVALID_WRITE_BACK_POLICY	Storage_Controller.Logical_Disk.INVALID_WRITE_BACK_POLICY	SIO06.02.0000002	3 (Error)	Check the Raid controller battery status, a fault in the battery will impact the Write-Back functionality.
Logical Disk access mode is not set to "Read-Write"	LOGICAL_DISK_INVALID_ACCESS_MODE	Storage_Controller.Logical_Disk.INVALID_ACCESS_MODE	SIO06.02.0000003	5 (Critical)	Access the storage controller and change the logical disk's access mode to "Read-Write".
The Logical Disk RAID level is not set to RAID0	LOGICAL_DISK_INVALID_RAID_LEVEL	Storage_Controller.Logical_Disk.INVALID_RAID_LEVEL	SIO06.02.0000004	3 (Error)	The RAID type is set in the storage controller is incorrect. Access the storage controller, verify that the Logical Disk does not contain any data that is not backed-up, destroy the Logical Disk,

Table 34 VxFlex Ready Node–specific Alerts in SNMP, VxFlex OS GUI, REST, and ESRS (continued)

Alert Message in VxFlex OS GUI	Alert Message in REST	Alert Message in SNMP Trap	Alert Code (for ESRS)	Severity	Recommended Action
					and re-create the logical disk as RAID0.
Logical Disk caching policy is not set to "DirectIO"	LOGICAL_DISK_INVALID_CACHE_POLICY	Storage_Controller.Logical_Disk.INVALID_CACHE_POLICY	SIO06.02.0000005	3 (Error)	Access the storage controller, and change the RAID0 caching policy for the Logical Disk to "DirectIO".
Logical disk is no longer being cached	LOGICAL_DISK_NO_LONGER_CACHED	Storage_Controller.Logical_Disk.LOGICAL_DISK_NO_LONGER_CACHED	SIO06.02.0000006	3 (Error)	Disable the Physical Drive physical cache. Note: This may cause DI issues in power cycle scenarios.
The state of the backup battery in the storage controller is not optimal	BATTERY_INVALID_STATE	Storage_Controller.Battery.Invalid_State	SIO06.03.0000001	5 (Critical) 3 (Error)	The backup battery is not fully charged, but it will recharge itself while the storage controller is powered on.
The backup battery in the storage controller needs to be replaced	BATTERY_REPLACEMENT_REQUIRED	Storage_Controller.Battery.REPLACEMENT_REQUIRED	SIO06.03.0000002	5 (Critical)	The backup battery in the storage controller may be nearing the end of its working life. Replace the battery.
The storage controller battery temperature is not within the configured threshold	BATTERY_INVALID_TEMPERATURE	Storage_Controller.Battery.INVALID_TEMPERATURE	SIO06.03.0000003	3 (Error)	Check the temperature in the server and at the data center. Check if a fan alert is raised in the node. Check the battery and the RAID controller and replace faulty items.
There is no backup battery installed in the storage controller	BATTERY_NOT_PRESENT	Storage_Controller.Battery.NOT_PRESENT	SIO06.03.0000004	5 (Critical)	Install a backup battery in the storage controller.
The storage controller contains an incompatible battery pack	BATTERY_INVALID_PACK_ENERGY	Storage_Controller.Battery.Invalid_Pack_Energy	SIO06.03.0000005	3 (Error)	Replace the storage controller battery with one that is compatible with your controller's model.
RAID controller's battery has invalid voltage state	BATTERY_INVALID_VOLTAGE	Storage_Controller.Battery.Invalid_Voltage	SIO06.03.0000006	3 (Error)	Check the RAID controller's battery. It may need to be replaced.
Hypervisor boot drive state is not OK	BOOT_DRIVE_INVALID_STATE	Node.Boot_Drive.Invalid_State	SIO07.01.0000001	3 (Error)	The drive on which the ESXi hypervisor is installed has reported an error. Replace the drive and reinstall the ESXi on the new drive.

Table 34 VxFlex Ready Node–specific Alerts in SNMP, VxFlex OS GUI, REST, and ESRS (continued)

Alert Message in VxFlex OS GUI	Alert Message in REST	Alert Message in SNMP Trap	Alert Code (for ESRS)	Severity	Recommended Action
The storage controller is not operating optimally	STORAGE_CONTROLLER_INVALID_STATE	Node.Storage_Controller.Invalid_State	SIO07.02.0000001	5 (Critical)	The storage controller may be faulty, and should be replaced.
The storage controller temperature has exceeded the configured threshold	STORAGE_CONTROLLER_INVALID_TEMPERATURE	Node.Storage_Controller.Invalid_Temperature	SIO07.02.0000002	5 (Critical) 3 (Error)	Ensure that the server and its environment are properly cooled and that all the chassis fans are functional.
CacheCade license is not installed	STORAGE_CONTROLLER_CACHECADE_NOT_LICENSED	Node.Storage_Controller.Cachecade_Not_Licensed	SIO07.02.0000003	3 (Error)	Install a CacheCade license on the storage controller.
Some of the disk slots in the storage controller are empty	STORAGE_CONTROLLER_NOT_ALL_SLOTS_FULL	Node.Storage_Controller.Not_All_Slots_Full	SIO07.02.0000004	3 (Error)	Add more disks if needed. If a disk was removed during a FRU (Field Replacement Unit) operation, insert the new disk into the chassis.
Unable to query RAID controller due to insufficient permissions	NODE_FAILED_TO_QUERY_RAID_CONTROLLER_INVALID_PERMISSIONS	Node.Storage_Controller.NODE_FAILED_TO_QUERY_RAID_CONTROLLER_INVALID_PERMISSIONS	SIO07.02.0000005	5 (Critical)	Verify that the user name and password of the Storage VM match the user name and password configured in the AMS (System Settings menu > System Settings > Security option). Resolve the differences and try again.
The physical CPU socket is not enabled	SOCKET_DISABLED	Node.Cpu_Socket.Disabled	SIO07.03.0000001	3 (Error)	Enable the CPU socket in the server BIOS.
Not all CPU cores are enabled	SOCKET_NOT_ALL_CORES_ENABLED	Node.Cpu_Socket.NOT_ALL_CORES_ENABLED	SIO07.03.0000002	3 (Error)	Change the physical CPU core configuration in the server BIOS.
CPU is not operating at full speed	SOCKET_SPEED_IS_NOT_MAX_SPEED	Node.Cpu_Socket.SPEED_IS_NOT_MAX_SPEED	SIO07.03.0000003	3 (Error)	Set the CPU clock speed in the server's BIOS to its maximum speed.
CPU temperature is not within threshold	CPU_SOCKET_INVALID_TEMPERATURE	Node.Cpu_Socket.INVALID_TEMPERATURE	SIO07.03.0000004	5 (Critical) 3 (Error)	The server's CPU temperature has exceeded the configured threshold. Make sure that the server is properly cooled and that the CPU and internal chassis fans are active.
The CPU Voltage Regulator's temperature is	CPU_SOCKET_INVALID_VR_TEMPERATURE	Node.Cpu_Socket.INVALID_VR_TEMPERATURE	SIO07.03.0000005	5 (Critical) 3 (Error)	Make sure that the server is properly cooled and that the

Table 34 VxFlex Ready Node–specific Alerts in SNMP, VxFlex OS GUI, REST, and ESRS (continued)

Alert Message in VxFlex OS GUI	Alert Message in REST	Alert Message in SNMP Trap	Alert Code (for ESRS)	Severity	Recommended Action
higher than the configured threshold					CPU and internal chassis fans are active.
The CPU Voltage Regulator's voltage is not within the configured threshold	CPU_SOCKET_INVALID_VR_VOLTAGE	Node.Cpu_Socket.INVALID_VR_VOLTAGE	SIO07.03.0000006	5 (Critical) 3 (Error)	1. Verify that the power supply is functioning correctly. 2. Try to replace a port in the Power Distribution Unit, or supply an external power source to check it. 3. Replace the power cable. 4. Replace the Power Supply Unit module.
CPU has invalid voltage state state	CPU_SOCKET_INVALID_VOLTAGE	Node.Cpu_Socket.Invalid_Voltage	SIO07.03.0000007	5 (Critical)	Check the chassis power supply
RAM temperature is not within threshold	RAM_INVALID_TEMPERATURE	Node.Ram.Invalid_Temperature	SIO07.04.0000001	5 (Critical) 3 (Error)	The temperature of one or more Server RAM modules exceeds the configured threshold. Ensure that the server and its environment are properly cooled and that all the chassis fans are functional.
RAM voltage regulator's temperature is not within threshold	RAM_INVALID_VR_TEMPERATURE	Node.Ram.INVALID_VR_TEMPERATURE	SIO07.04.0000002	5 (Critical) 3 (Error)	Ensure that the server and its environment are properly cooled and that all the chassis fans are functional.
RAM voltage regulator's voltage is not within threshold	RAM_INVALID_VR_VOLTAGE	Node.Ram.INVALID_VR_VOLTAGE	SIO07.04.0000003	5 (Critical) 3 (Error)	Ensure that the server and its environment are properly cooled and that all the chassis fans are functional. If the system still issues an alert, the DIMM may be faulty, and may have to be replaced.
The node temperature has exceeded the configured threshold	NODE_INVALID_TEMPERATURE	Node.Node.INVALID_TEMPERATURE	SIO07.05.0000001	5 (Critical) 3 (Error)	Ensure that the server and its environment are properly cooled and that all the chassis fans are functional.
The voltage values on the node are not within the configured threshold	NODE_INVALID_VOLTAGE	Node.Node.INVALID_VOLTAGE	SIO07.05.0000002	5 (Critical) 3 (Error)	1. Verify that the power supply is functioning correctly, and then try to replace a port in the Power Distribution Unit or check with external power. 2.

Table 34 VxFlex Ready Node–specific Alerts in SNMP, VxFlex OS GUI, REST, and ESRS (continued)

Alert Message in VxFlex OS GUI	Alert Message in REST	Alert Message in SNMP Trap	Alert Code (for ESRS)	Severity	Recommended Action
					Replace the power cable. 3. Replace the Power Supply Module.
Unable to connect to monitoring agent (BMC)	NODE_FAILED_TO_CONNECT_TO_BMC	Node.Node.FAILED_TO_CONNECT_TO_BMC	SIO07.05.0000004	5 (Critical)	The BMC IP address or User credentials do not respond to AMS queries. Verify network connectivity from the AMS to the BMC IP address, and check that the BMC admin user and password have not been tampered with. Contact customer support for assistance if the problem persists.
SDC is not installed on this node	NODE_WITH_NO_SDC	Node.Node.NODE_WITH_NO_SDC	SIO07.05.0000005	3 (Error)	Consider installing an SDC on this node, so that it can use VxFlex OS volumes
Unable to connect to monitoring agent (SVM)	NODE_FAILED_TO_CONNECT_TO_SVM	Node.Node.FAILED_TO_CONNECT_TO_SVM	SIO07.05.0000006	5 (Critical)	The SVM IP address or User credentials do not respond to AMS queries. Check the SVM state in vCenter and try to open the SSH or the console to it. Then check the message log for errors.
Unable to connect to monitoring agent (ESX)	NODE_FAILED_TO_CONNECT_TO_ESX	Node.Node.FAILED_TO_CONNECT_TO_ESX	SIO07.05.0000007	5 (Critical)	The ESX IP address or User credentials do not respond to AMS queries. Check the ESX state in vCenter and try to open the SSH or the console (via BMC http) to it. Then check the messages log for errors.
Unable to connect to monitoring agent (host)	NODE_FAILED_TO_CONNECT_TO_HOST	Node.Node.FAILED_TO_CONNECT_TO_HOST	SIO07.05.0000007	5 (Critical)	The ESX IP address or User credentials do not respond to AMS queries. Check the ESX state in vCenter and try to open the SSH or the console (via BMC http) to it. Then check the messages log for errors.
Unable to connect to monitoring agent (VCenter)	NODE_FAILED_TO_CONNECT_TO_VCENTER	Node.Node.FAILED_TO_CONNECT_TO VCENTER	SIO07.05.0000008	5 (Critical)	Ensure that the vCenter configuration is correct, and ensure that the node's Mgmt port is routable to the vCenter IP address

Table 34 VxFlex Ready Node–specific Alerts in SNMP, VxFlex OS GUI, REST, and ESRS (continued)

Alert Message in VxFlex OS GUI	Alert Message in REST	Alert Message in SNMP Trap	Alert Code (for ESRS)	Severity	Recommended Action
Node serial number has changed	NODE_SERIAL_NUMBER_CHANGE	Node.Node.SERIAL_NUMBER_CHANGE	SIO07.05.0000009	3 (Error)	The S/N of the Motherboard does not match the IP address of the ESX. Either the SATADOM was moved to a new server or the motherboard was replaced. It is also possible that a new server that has the same IP addresses and the same user and password, but was not installed using the AMS flow, exists in the network. EMC does not support such field replacement cases. Such cases should be solved by replacing nodes using EMC-recommended procedures. For more information, contact Customer Support.
RFcache pool state is "Not Started"	XTREMCACHE_INVALID_STATE	Node.Node.XTREMCACHE_INVALID_STATE	SIO07.05.0000010	3 (Error)	Check the Read Flash Cache pool state. The cache device might be misconfigured. In this case, remove the cache device from the cache pool and add it back again.
Host certificate is about to expire	HOST_CERTIFICATE_ABOUT_TO_EXPIRE	Node.Node.HOST_CERTIFICATE_ABOUT_TO_EXPIRE	SIO07.05.0000014	3 (Error)	Renew the AMS certificate and then run "renew certificate" process
CMOS battery state is invalid	NODE_INVALID_CMOS_BATTERY	Node.Node.Invalid_Cmos_Battery	SIO07.05.0000015	3 (Error)	Check the node's CMOS battery. It may need to be replaced.
The node is not being managed by the current AMS.	NODE_NOT_MANAGED_BY_AMS	Node.Node.Node_Not_Managed_By_Ams	SIO07.05.0000016	5 (Critical)	Verify that the node is not being managed by a different AMS.
Failed to connect to the SVM due to authentication failure	NODE_FAILED_TO_CONNECT_TO_SVM_AUTH_FAILED	Node.Node.Node_Failed_To_Connect_To_Svm_Auth_Failed	SIO07.05.0000017	5 (Critical)	Change the SVM's root password to be the one set by the AMS
Failed to connect to the host. The SSH key has been changed, and is not trusted by the AMS.	NODE_FAILED_TO_CONNECT_TO_HOST_SSH_KEY_MISMATCH	Node.Node.Node_Failed_To_Connect_To_Host_Ssh_Key_Mismatch	SIO07.05.0000018	5 (Critical)	This is a security issue. Remove the node, and add the node again to the AMS.

Table 34 VxFlex Ready Node–specific Alerts in SNMP, VxFlex OS GUI, REST, and ESRS (continued)

Alert Message in VxFlex OS GUI	Alert Message in REST	Alert Message in SNMP Trap	Alert Code (for ESRS)	Severity	Recommended Action
Failed to connect to the SVM. The SSH key has been changed, and is not trusted by the AMS.	NODE_FAILED_TO_CONNECT_TO_SVM_SSH_KEY_MISMATCH	Node.Node.Node_Failed_To_Connect_To_Svm_Ssh_Key_Mismatch	SIO07.05.0000019	5 (Critical)	This is a security issue. Remove the node, and add the node again to the AMS.
Failed to connect to the host because the certificate is not trusted	NODE_FAILED_TO_CONNECT_TO_ESX_CERTIFICATE_NOT_TRUSTED	Node.Node.Node_Failed_To_Connect_To_Esx_Certificate_Not_Trusted	SIO07.05.0000020	5 (Critical)	This is a security issue. Remove the node, and add the node again to the AMS.
Failed to connect to the host due to authentication failure	NODE_FAILED_TO_CONNECT_TO_HOST_AUTH_FAILED	Node.Node.Node_Failed_To_Connect_To_Host_Auth_Failed	SIO07.05.0000021	5 (Critical)	Change the host's root password to be the one set by the AMS
The MDM has a self-signed certificate that was not replaced by the AMS. Connection to the MDM is not authenticated by the AMS.	AMS_MDM_HAS_SELF_SIGNED_CERTIFICATE	Node.Node.Ams_Mdm_Has_Self_signed_certificate	SIO07.05.0000022	5 (Critical)	Disable AMS secure connection, and enable it again in order to automatically sign the MDM certificate. If the previous solution does not work, remove and add the the MDM again in secure mode.
CPU cache is not enabled	SOCKET_CACHE_DISABLED	Cpu_Socket.Socket_Cache.Disabled	SIO08.01.0000001	3 (Error)	Enable the CPU cache in the server BIOS.
CPU cache size is not optimal	SOCKET_CACHE_SIZE_NOT_MAX_SIZE	Cpu_Socket.Socket_Cache.SIZE_NOT_MAX_SIZE	SIO08.01.0000002	3 (Error)	Check the CPU cache in the server BIOS, and set it to maximum size.
Fan speed is not within threshold	FAN_INVALID_SPEED	Chassis.Fan.Invalid_Speed	SIO09.01.0000001	5 (Critical) 3 (Error)	Ensure that the server is properly cooled and that the chassis fans are functional.
The power supply input wattage is not within the configured threshold	PSU_INVALID_INPUT	Chassis.Psu.Invalid_Input	SIO09.02.0000001	5 (Critical) 3 (Error)	1. Verify that the power supply is functioning correctly. 2. Try to replace a port in the Power Distribution Unit, or supply an external power source to check. 3. Replace the power cable. 4. Replace the Power Supply Unit module.
Power supply is missing	PSU_NOT_AVAILABLE	Chassis.Psu.Not_Available	SIO09.02.0000002	3 (Error)	Install a new Power Supply Unit, or if there is an existing PSU, verify that it is properly connected.
Power supply power loss	PSU_POWER_LOST	Chassis.Psu.Power_Lost	SIO09.02.0000003	5 (Error)	Check the power outlet/ power feed to the server.

Table 34 VxFlex Ready Node–specific Alerts in SNMP, VxFlex OS GUI, REST, and ESRS (continued)

Alert Message in VxFlex OS GUI	Alert Message in REST	Alert Message in SNMP Trap	Alert Code (for ESRS)	Severity	Recommended Action
Network Adapter Link is down	NIC_PORT_DOWN	Nic.Nic_Port.Nic_Port_Down	SIO11.01.0000001	5 (Critical)	The network link is down in the network adapter. Verify that the network adapter is enabled in the operating system, and verify that the network cable is connected properly at both ends.
Drive is ready to be removed	PHYSICAL_DRIVE_REMOVED_FROM_OS	Node.Physical_Drive.Ready_For_Removal	SIO07.06.0000001	2(Warning)	The drive is ready to be removed from the node.
BOSS is not configured with expected Raid Level 1	STORAGE_CONTROLLER_INVALID_BOSS_RAID_LEVEL	Node.Boot_Drive.Invalid_Boss_Raid_Level	SIO07.01.0000002	5 (Critical)	Check the BOSS setup configuration
Interleave must be enabled in BIOS when NVDIMM devices are present	NVDIMM_INTERLEAVE_DISABLED	Node.Ram.Invalid_Interleave_Bios_setting	SIO07.04.0000004	5 (Critical)	Enable NVDIMM interleave in BIOS settings.
Persistent Memory must be enabled in BIOS when NVDIMM devices are present	NVDIMM_PERSISTENT_MEMORY_DISABLED	Node.Ram.Invalid_Persistent_Memory_Bios_setting	SIO07.04.0000005	5 (Critical)	Enable Persistent Memory in BIOS settings
NVDIMM battery state is low.	NVDIMM_BATTERY_INVALID_STATE	Node.Ram.NVDIMM_Battery.Invalid_State	SIO07.04.0000006	5 (Critical) 3 (Error)	Check the NVDIMM battery status. A fault in the battery will impact the NVDIMM persistent memory. Persistent memory will not function if battery is faulty.
Install or check GPU drivers	GPUCARD_MONITORING_ERROR	Node.Gpu.Monitoring_Error	SIO07.06.0000001	2(Warning)	Check that the GPU drivers are properly installed on the node.
Replace DIMM	DIMM_IN_ERROR_STATE	Node.Ram.Error_State	SIO07.04.0000007	3 (Error)	Replace DIMM
DIMM reports Degraded state	DIMM_IN_DEGRADED_STATE	Node.Ram.Degraded_State	SIO07.04.0000008	2(Warning)	Replace DIMM as soon as possible

Configure SNMP properties after deployment

SNMP properties can be configured using the VxFlex OS GUI.

The following procedures are required to enable the SNMP feature:

1. Creating a Lockbox

2. Configuring SNMP using the VxFlex OS GUI

Create a lockbox

Create a lockbox and add the MDM credentials. Lockbox is required for the following features: ESRS, SNMP, LDAPS.

Before you begin

The following items are required for the ESRS feature:

- ESRS gateway v3 version 3.08 or higher must be installed and configured. It is recommended to create at least two ESRS Gateways and to define them as a cluster via the backend server.
- ESRS gateway must be reachable from VxFlex OS on port 9443.
- The VxFlex OS license must be installed.

Ensure you have:

- One or more IP addresses of the ESRS gateway servers. Note that ESRS does not currently support IPv6.
- ESRS username and password.
- VxFlex OS Gateway IP address, username, and password.
- MDM username and password.
- The VxFlex OS Management IP address to be used as the Connect-In IP address. It must be an IP address that is accessible from the ESRS gateway (for example, in case of NAT).

Use FOSGWTool to configure a Lockbox. FOSGWTool should be used to create a Lockbox only when a Lockbox has not yet been created.

To use FOSGWTool, input the appropriate path, based on your operating system, and append the commands to the end of the filepath:

- Linux FOSGWTool filepath: `/opt/emc/scaleio/gateway/bin/FOSGWTool.sh`
- Windows FOSGWTool filepath: `C:\Program Files\EMC\ScaleIO\Gateway\bin\FOSGWTool.bat`

Procedure

1. Create a Lockbox:

```
<FOSGWTool_PATH> --change_lb_passphrase --new_passphrase
<NEW_PASSPHRASE>
```

Note

From system version 2.5 and later, the installation process will assign a random passphrase to this property, and it is highly recommended not to configure or use this property, because it could create a security breach.

Windows example:

```
 C:\Program Files\EMC\ScaleIO\Gateway\bin\FOSGWTool.bat --
set_mdm_credentials --mdm_user admin --mdm_password Scaleio123
```

2. Add MDM credentials to the Lockbox:

```
<FOSGWTool_PATH> --set_mdm_credentials --mdm_user
<MDM_USERNAME> --mdm_password <MDM_PASSWORD>
```

Configuring SNMP using the VxFlex OS GUI

Configure SNMP and add trap receivers using the VxFlex OS GUI.

Before you begin

- Ensure that you can authenticate to the VxFlex OS GUI with an admin user.
- The MDM native user/LDAP user in the lockbox must be assigned Monitor privileges in addition to Backend Configurator privileges. The Backend Configurator role does not automatically include Monitor privileges.

For more information on SNMP use in VxFlex OS, see the user documentation.

Procedure

1. Log in to the VxFlex OS GUI.
2. From the **System Settings** menu at the top right side of the window, in any view, select **System Settings** > **SNMP**.
3. In the **Set SNMP Configuration** panel, add the first Trap Receiver by typing its IP address in the **Trap Receiver** box, and the port number in the **Port** box.
4. To add another Trap Receiver, click the **Add** button . In the new row, type the additional IP address in the new **Trap Receiver** box, and the port number in the new **Port** box.
5. Type a number in the **Resend Frequency** box.

   Typing "0" will make the system use the default, which is a resend frequency of 10 seconds.
6. To send a test SNMP trap, click **Test**.
7. Click **Save**.

Results

The SNMP configuration has been added to the system.

VxFlex OS.mib file

The following text is the content of the `scaleio.mib` file.

Note

The object source identifier in the trap MDM.MDM_Cluster.SNMP_Server_Cannot_Connect_To_MDM is "NA".

```
SCALEIO-MIB DEFINITIONS ::= BEGIN IMPORTS MODULE-IDENTITY,
OBJECT-TYPE, NOTIFICATION-TYPE, Integer32 FROM SNMPv2-SMI
```

```
DisplayString FROM RFC1213-MIB OBJECT-GROUP, NOTIFICATION-GROUP
FROM SNMPv2-CONF emc FROM EMC-MIB;
```

```
scaleio MODULE-IDENTITY
LAST-UPDATED "201511060000Z"
ORGANIZATION "EMC Corporation"
CONTACT-INFO
"EMC Corporation

www.emc.com"

DESCRIPTION
"The Structure of Management Information for the EMC SCALEIO
enterprise."
REVISION      "201511060000Z"
DESCRIPTION
"Initial version of this MIB."
::= { emc 101 }
```

```
-- 1.3.6.1.4.1.1139.101.1
scaleioAlert OBJECT IDENTIFIER ::= { scaleio 1 }
```

```
-- 1.3.6.1.4.1.1139.101.1.1
scaleioAlertSeverity OBJECT-TYPE
SYNTAX Integer32
MAX-ACCESS accessible-for-notify
STATUS current
DESCRIPTION
"Severity of the event"
::= { scaleioAlert 1 }
```

```
-- 1.3.6.1.4.1.1139.101.1.2
scaleioAlertType OBJECT-TYPE
SYNTAX DisplayString (SIZE (0..255))
MAX-ACCESS accessible-for-notify
STATUS current
DESCRIPTION
"Type of the alert"
::= { scaleioAlert 2 }
```

```
-- 1.3.6.1.4.1.1139.101.1.3
scaleioAlertSourceObjectId OBJECT-TYPE
SYNTAX DisplayString (SIZE (0..255))
MAX-ACCESS accessible-for-notify
STATUS current
DESCRIPTION
"Object id for which the alert was created"
::= { scaleioAlert 3 }
```

```
-- 1.3.6.1.4.1.1139.101.1.4
scaleioAlertActionCode OBJECT-TYPE
SYNTAX DisplayString (SIZE (0..255))
MAX-ACCESS accessible-for-notify
STATUS current
DESCRIPTION
```

```
"Action code of the alert"
::= { scaleioAlert 4 }
```

```
-- 1.3.6.1.4.1.1139.101.1.5
scaleioGroups OBJECT IDENTIFIER ::= { scaleioAlert 5 }
```

```
-- 1.3.6.1.4.1.1139.101.1.5.1
currentObjectGroup OBJECT-GROUP
OBJECTS { scaleioAlertSeverity,
scaleioAlertType,
scaleioAlertSourceObjectId,
scaleioAlertActionCode }
STATUS current
DESCRIPTION
"scaleio-MIB-V2 OBJECT-GROUP."
::= { scaleioGroups 1 }
```

```
-- 1.3.6.1.4.1.1139.101.1.5.2
currentNotificationGroup NOTIFICATION-GROUP
NOTIFICATIONS { scaleioAEAlert }
STATUS current
DESCRIPTION
"scaleio-MIB-V2 NOTIFICATION-GROUP."
::= { scaleioGroups 2 }
```

```
scaleioAEAlert NOTIFICATION-TYPE
OBJECTS { scaleioAlertSeverity, scaleioAlertType,
scaleioAlertSourceObjectId, scaleioAlertActionCode }
STATUS current
DESCRIPTION "ScaleIO Alert"
::= { scaleio 2 }
END
```

CHAPTER 14

Server Factory Reset and Power Operations

The following topics explain how to open a Dell iDRAC console session, how to perform power operations, and how to perform a factory reset.

- Perform a factory reset 314
- KVM activities 322

Perform a factory reset

Return a server to the factory-installed state, including installation of the operating system, preceded by setting the BIOS settings.

Before you begin

Ensure that you have the specific ISO file supplied in your VxFlex OS software package.

The following list is a summary of the procedures that must be performed in order to return a server to factory default settings.

Procedure

1. Set up the BMC (iDRAC) IP address and validate the BIOS, as described in "Set up the BMC (IDRAC) IP address and BIOS".
2. Verify the status of system hardware, storage controller, and disks, as described in "Verify the status of the system hardware, storage controller, and disks".
3. Open the iDRAC console, as described in"Open the KVM console".
4. Update the BIOS, firmware and settings, as described in "Updating the BIOS, firmware and settings".
5. Map the ISO file, as described in "Map the ISO file on a server."

Set up the BMC (iDRAC) IP address and BIOS

Set up the BMC (iDRAC) IP address and set up or validate the BIOS on the VxFlex Ready Node servers.

Before you begin

Ensure that you have access to, or have the details for:

- The KVM console
- The server BMC (iDRAC) IP address
- The server BMC (iDRAC) subnet mask
- The gateway IP address
- The VLAN ID of the BMC (iDRAC), if the VLAN is used

Note

The BMC (iDRAC) IP address may be an IPv4 or an IPv6 address.

Defaults:

- BMC (iDRAC) username - root
- BMC (iDRAC) password - The default password is calvin for 13G nodes and Scaleio123 for 14G nodes.
- BIOS password - emcbios

During the BMC (iDRAC) IP and BIOS setup, use the following operations:

- Use the arrow keys to navigate in the BIOS screens.
- Use the + and - keys on the keyboard to change the option selection in the BIOS screens.

- Use SPACE or ENTER keys to change the settings.
- The ENTER key opens a list to the desired values.

Procedure

1. Open the KVM console.

   For console operations (KVM access), ensure that you have either a VGA tool kit/Crash Cart to allow physical console connection from a laptop computer to a server, or a computer screen and keyboard connection to the rack.

2. Press F2 to access the main menu.
3. From the **System Setup Main Menu** screen, select the **iDRAC Settings** menu option.
4. Configure network settings:

   a. In the **iDRAC Settings** screen, select **Network**.

   b. In the **iDRAC Settings--Network** pane, verify the following parameter values:

      - **Enable NIC** = **Enabled**
      - **NIC Selection** = **Dedicated**

   c. From the **IPv4 Settings** pane, configure the IPv4 parameter values for the BMC (iDRAC) port:

      - **Enable IP IPv4** = **Enabled**
      - **Enable DHCP** = **Disabled**
      - **Static IP Address** = Static IP address
      - **Static Gateway** = Gateway IP address
      - **Static Subnet Mask** = Subnet mask IP address

   d. From the **IPv6 Settings** pane, configure the IPv6 parameter values for the BMC (iDRAC) port.

   e. From the **IPMI Settings** pane, verify the following parameter values:

      - **Enable IPMI Over LAN** = **Enabled**
      - **Channel Privilege Level Limit** = **Administrator**

   f. If you are working in a VLAN setup, access the **VLAN Configuration** pane and configure the VLAN ID parameters.

   g. When the parameter set up is complete, click **Back** to display the **iDRAC Settings** screen.

5. From the **iDRAC Settings** screen, click **Finish**, **Yes**, and then **OK** to return to the **System Setup Main Menu** screen.
6. In the **System Setup Main Menu** screen, select the **System BIOS** menu.
7. In the **System BIOS Settings** screen, verify that the processor settings are correct.

   If the settings are incorrect, configure them as follows:

   a. Select **Processor Settings**.

   b. In the **System BIOS Settings--Processor Settings** pane, verify the following parameter values:

- **Virtualization Technology** = **Enabled**
- **Number of Cores Per Processor** = **All**

c. Click **Back** to return to the **System BIOS Settings** screen.

8. Configure boot settings,:
   a. Select **Boot Settings**.
   b. In the **System Boot Settings--Boot Settings** pane, verify that **Boot Mode** is set to **BIOS (13G) or UEFI (14G)**.
   c. Select the **BIOS Boot Settings** link.
   d. In the **System BIOS Settings--Boot Settings--BIOS Boot Settings** pane, verify that:
      - In the **Boot Sequence** list, **Hard drive C:** appears as the first item.
      - In the **Hard-Disk Drive Sequence** list, the **SATADOM-ML 3SE (13G) or BOSS (14G)** device appears as the first item.
   e. Click **Back** twice to return to the **System BIOS Settings** screen.
9. Configure integrated devices:
   a. Select **Integrated Devices**.

      The **System BIOS Settings--Integrated Devices** screen appears.
   b. Set **SR-IOV Global Enable** to **Enabled**.
   c. Verify that the **Internal USB Port** parameter is set to **Off**.
   d. Click **Back** to return to the **System BIOS Settings** screen.
10. From the **System BIOS Settings** screen, click **Finish**, **Yes**, and then **OK** to return to the **System Setup Main Menu** screen.
11. Select **Finish** to exit the BIOS and apply all settings post boot.

Results

The BMC (iDRAC) IP and server BIOS address configuration is complete.

Verify the status of the system hardware, storage controller, and disks - 13G servers

Use the following procedure to verify the status of the system hardware, storage controller, and disks in a VxFlex Ready Node 13G server.

Before you begin

Ensure that you know:

- The IP address of the BMC (iDRAC) port
- The username and password for the BMC (iDRAC) portal (default username and password are root and password)

Procedure

1. From a browser, go to http://*<BMC/iDRAC_IP_address>*.
2. In the **DELL Console Login** window, type the user name and password, and then click **Login**.

The **System Summary** dashboard displays the high-level status of all hardware devices in the **Server Health** pane.
In an ideal scenario, all hardware sensors are green.

3. In the **Quick Launch Tasks** pane, select **View Logs**.

   The **System Event Log** is displayed with color-coded severity levels.

4. Ensure that any power supply- and fan-related events in the event log are non-repetitive. Repetitive events may be due to the intermittent nature of faults, such as poor physical connections.

   For repetitive events, it is recommended that you remove the relevant hardware module and replace it in its socket.

5. In the navigation pane, select **Storage** > **Controllers**.

   The controller-related information is displayed in the **Health and Properties** table, with the controller type shown in the **Name** column.

   For example:

   ```
   PERC H730 Mini (Embedded)
   ```

6. In the navigation pane, select **Storage** > **Physical Disks**.

   Information about the physical disks is displayed in a table in the **Health and Properties** pane.

7. Verify that no disk is in failed state.

   If any of the disks is failed, refer to the relevant Disk Replace FRU.

8. In the navigation pane, select **Server**.
9. In the **System Inventory** pane, verify that the server drivers and firmware in the **Firmware Inventory** list match the required versions, as published in the Driver and Firmware Matrix.

   If the driver and firmware versions do not match the matrix, the AMS deployment will align the drivers and firmware with the latest supported versions.

Verify the status of the system hardware, storage controller, and disks - 14G servers

Verify the status of the system hardware, storage controller, and disks in a VxFlex Ready Node 14G server.

Before you begin

Ensure that you know:

- The IP address of the BMC (iDRAC) port
- The username and password for the BMC (iDRAC) portal (default username and password are root and password)

Procedure

1. From a browser, go to http://*<BMC/iDRAC_IP_address>*.
2. In the **DELL Console Login** window, type the user name and password, and then click **Login**.

   The **Dashboard** displays the high-level status of all hardware devices in the **System Health** pane.

In an ideal scenario, all hardware sensors are green.

3. Navigate to **Maintenance** > **System event log.**

   The **System Event Log** pane is displayed with color-coded severity levels.

4. Ensure that any power supply- and fan-related events in the event log are non-repetitive. Repetitive events may be due to the intermittent nature of faults, such as poor physical connections.

   For repetitive events, it is recommended that you remove the relevant hardware module and replace it in its socket.

5. In the navigation pane, select **Storage** > **Overview** > **Controller**.

   A table displays the controller-related information, with the controller type in the **Name** column. For example:

   ```
   PERC H740P Mini (Embedded)
   ```

6. In the navigation pane, select **Storage** > **Physical Disks**.

   A table displays information regarding the physical disks, and an option to **Blink/Unblink** the selected disk.

7. Verify that no disk is in failed state.

   If any of the disks is failed, refer to the relevant Disk Replace FRU.

8. In the navigation pane, select **System**.

9. In the **Inventory** pane, verify that the server drivers and firmware in the **Firmware Inventory** list match the required versions, as published in the Driver and Firmware Matrix.

   If the driver and firmware versions do not match the matrix, the AMS deployment will align the drivers and firmware with the latest supported versions.

Open the KVM console - 13G servers

Open the KVM console on a VxFlex Ready Node 13G server.

Before you begin

Ensure that:

- The system environment meets the prerequisites for using the KVM console.
- You know the IP address of the BMC (iDRAC) port.
- You know the username and password for accessing the BMC (iDRAC) (default username/password are root/password).

Procedure

1. From your Internet browser, go to https://*<iDRAC_IP_address>*.
2. In the **DELL Console Login** window, type the user name and password, and click **Login**.
3. In the main KVM console window, select the **Server** node in the navigation pane.

   The **System Summary** dashboard is displayed.

4. In the **Virtual Console Preview** pane, click **Launch** to start a console session.

   A popup security warning screen is displayed.

5. Select **Accept** and click **Run**.

   The **Java Console** window opens and provides you with console access to the server.
   If this is the first time that you are opening a console, additional warning and confirmation prompts may appear. Click **OK** to grant approvals in these prompts.

   The subsequent steps depend on your browser selection and how it is configured. If downloads run automatically, the console window appears. If not, follow the instructions given in the next step to open the console window.

6. Depending on your browser, perform the following steps:

Browser	Steps
Firefox	a. Click **Launch Virtual Console**. The **"What should Firefox do with this file?"** window appears, with the **Open with Java(™) Web Start Launcher** option selected. b. Click **OK**. c. Scroll through the successive pop-up windows, then click **Run** to launch the console.
Google Chrome	a. Click **Show all downloads**. b. In the **Downloads** window, at the warning message, click **Keep**. c. In the **Downloads** list, click the downloaded file URL. d. At the warning message, click **Keep** or **Keep anyway**. e. Click the `jviewer.jnlp` file. If a security message appears, click **Run**. The management console window is displayed.
Internet Explorer	Click **Open** to run the `jviewer.jnlp` file.

Results

The KVM console is open and ready for use.

Open the KVM console - 14G servers

Open the KVM console on a VxFlex Ready Node 14G server.

Before you begin

Ensure that:

- The system environment meets the prerequisites for using the KVM console.
- You know the IP address of the BMC (iDRAC) port.
- You know the username and password for accessing the BMC (iDRAC) (default username/password are root/password).

Procedure

1. From your Internet browser, go to https://*<iDRAC_IP_address>*.
2. In the **DELL Console Login** window, type the user name and password, and click **Login**.

3. From the dashboard, click **Launch Virtual Console** to start a console session.

   A popup security warning screen is displayed.

4. Select **Accept** and click **Run**.

   The **Java Console** window opens and provides you with console access to the server.
   If this is the first time that you are opening a console, additional warning and confirmation prompts may appear. Click **OK** to grant approvals in these prompts.

   The subsequent steps depend on your browser selection and how it is configured. If downloads run automatically, the console window appears. If not, follow the instructions given in the next step to open the console window.

5. Depending on your browser, perform the following steps:

Browser	Steps
Firefox	a. Click **Launch Virtual Console**. The **"What should Firefox do with this file?"** window appears, with the **Open with Java(™) Web Start Launcher** option selected. b. Click **OK**. c. Scroll through the successive pop-up windows, then click **Run** to launch the console.
Google Chrome	a. Click **Show all downloads**. b. In the **Downloads** window, at the warning message, click **Keep**. c. In the downloads list, click the downloaded file URL. d. At the warning message, click **Keep** or **Keep anyway**. e. Click the `jviewer.jnlp` file. If a security message appears, click **Run**. The management console window is displayed.
Internet Explorer	Click **Open** to run the `jviewer.jnlp` file.

Results

The KVM console is open and ready for use.

Map the ISO file on a VxFlex Ready Node server

Map the ISO file on a VxFlex Ready Node server.

Before you begin

Ensure that you can access the console.

Procedure

1. If you have not already done so, open the KVM console.
2. In the **DELL System Setup** screen, select **Virtual Media** > **Connect Virtual Media**.

   The **Initializing connection** screen is displayed with a connecting message. You may have to wait for some time for the connecting message to finish.

3. In the **DELL System Setup** screen, select **Virtual Media** > **Map CD/DVD**.
4. From the **Virtual Media - Map CD/DVD** screen, browse and set up the VxFlex OS ISO file and then click **Map Device**.
5. To verify the VxFlex OS AMS Node ISO file selection (part of the AMS download page, located at https://support.emc.com/products/42215), click **Virtual Media** and view the ISO selection in the drop down list.

   You can select multiple supported ISO images: ESXi 6.0 P6, ESXi 6.5 P2, RHEL 7.3, or RHEL 7.4. The ISOs are patched with Specter and Meltdown fixes.

   Note

   For system board and BOSS card replacements, you can use the DTK ISO for initial firmware update. For information, see "DTK - Hardware Update Bootable ISO" in the reference section of this publication.

6. From the **Next Boot** menu, choose the Virtual CD/DVD/ISO.

   This option will make the server install the ISO the next time that the server boots up.
7. To reset the server, in the main **DELL System Setup** console window, select **Power** > **Reset System (warm boot)** to display the list of keyboard keys with related functionality.

Results

The server boots, using the OS ISO file, and finishes the installation. The ISO installation is automatic and will reboot the node twice during the process. No interaction is required. After the installation, the node will boot up to the operating system. For ESXi hosts, the SVM is deployed in listening mode, ready for an AMS detection.

Note

If you used the DTK ISO, for the next step refer to the "DTK - Hardware Update Bootable ISO" in the reference section of this publication.

Map the ISO file on a VxFlex Ready Node 14G server

Map the ISO file on a VxFlex Ready Node 14G server.

Before you begin

Ensure that you can access the console.

Procedure

1. Open the KVM console, using the **Launch** link.
2. In the **DELL System Setup** screen, select **Virtual Media** > **Connect Virtual Media**.

   The **Initializing connection** screen is displayed with a connecting message. You may have to wait for some time for the connecting message to finish.
3. In the **DELL System Setup** screen, select **Virtual Media** > **Map CD/DVD**.
4. From the **Virtual Media - Map CD/DVD** screen, browse and set up the `ISO` file and then click **Map Device**.
5. To verify the `ISO` file selection, click **Virtual Media** and view the ISO selection in the drop down list.

6. Reset the server:
   a. In the main **DELL System Setup** console window, select **Power** > **Reset System (warm boot)** to display the list of keyboard keys with related functionality.
   b. Press F11. The **Boot Manager Main Menu** pane is displayed.
   c. Select **One-shot BIOS Boot Menu**. The **Select Legacy Boot Option** pane is displayed.
   d. Select **Virtual Optical Drive**.
   a. In the main **DELL System Setup** console window, select **Power** > **Rest System (warm boot)** to display the list of keyboard keys with related functionality.
   b. Refer to the best practices of DELL OS installation: How to Install an Operating System on Dell PowerEdge Servers.

      Ensure that you select the BOSS device as the OS install target drive.

   The server boots using the OS ISO file and finishes the installation.
7. If the node contains a factory-installed OS, the datastores will not be visible in the vSphere Web Client. The BOSS card is partitioned, and the OS is installed in the visible partition. To use the remaining storage on the BOSS card for the datastore, perform the following steps:
   a. In the vSphere Web Client, select **Storage** > **New Datastore** > **Create New VMFS datastore**.
   b. Select the empty partition for the datastore.

      Datastore 1 is created on the empty BOSS card.

After you finish

Note

A Red Hat license is required for storage-only configurations with Red Hat OS image downloaded from the VxFlex OS support page.
Customers can choose a Red Hat licensing option listed on this page, or bring their own Red Hat license. Dell EMC is not responsible for enforcing OS licensing.

KVM activities

The KVM console allows you to perform additional server maintenance tasks.

Additional tasks using the KVM console - 13G server

You can perform the following additional tasks on a VxFlex Ready Node 13G server using the KVM console server.

Before you begin

Ensure that you can access the KVM console.

Procedure

1. Open the KVM console, as described in Open the KVM console - 13G servers on page 318.

2. You can perform the following tasks from the **System Summary** dashboard:

Task Option	Task Steps
Blink the server LED	a. In the **Quick Launch Tasks** pane, click **System ID LED ON/OFF** and then click **OK**. The LED starts blinking. b. To stop the LED blinking, click **System ID LED ON/OFF** again and then click **OK**. The LED stops blinking.
Retrieve the chassis serial number	a. Click the **Properties** menu option in the menu bar, select **Details**. The **System Details** window is displayed. b. In the **Main System Chassis** pane, check the chassis serial number in the **Node Id** field.
Power up/down the server	a. In the **Quick Launch Tasks** pane, click **Power ON/OFF**. b. Click **OK**.

Additional tasks using the KVM console for 14G servers

You can perform the following tasks from the KVM console server on a VxFlex Ready Node 14G server.

Before you begin

Ensure that you can access the KVM console.

Procedure

1. To activate the server LED, use one of the following methods:
   - On the dashboard, click **Identify System**, and then in the **Success** message click **OK**.
     The LED starts blinking.

     To stop the LED from blinking, click **Identify System** again, and then click **OK**.

     The LED stops blinking.
   - Display the front panel status LEDs:
     a. On the dashboard, select **Configuration** > **System settings** > **Hardware settings**.
     b. From the **Front Panel Configuration** menu, select **Monitor Server LED - From Dashboard** > **System** > **Overview**.
     c. Click **Front Panel**.
     The front panel status LEDs are displayed.
2. To retrieve the chassis serial number, in the dashboard click **System** > **Details**.

   The service tag is displayed with the corresponding serial number.
3. To power the server off or on, use of the following methods:
   - From the dashboard, click **Graceful Shutdown** > **Power Off**, and then click **OK**.

- From the dashboard, select **Power Management** > **Power Control** > **Power Off System / Power On System**.

CHAPTER 15

DTK - Hardware Update Bootable ISO

The Dell OpenManage Deployment Toolkit (DTK) includes a set of utilities, sample scripts, and sample configuration files that you can use to deploy and configure the Dell systems. This section provides additional information on using the VxFlex Ready Node Hardware Update Bootable ISO ("Hardware ISO") to update drivers, BIOS, and firmware on VxFlex Ready Node servers.

- Dell OpenManage DRAC Tools (RACADM) 326
- Update the SATADOM firmware (13G servers only) 328
- Recommended BIOS and firmware settings 329
- Troubleshooting the Hardware ISO 331
- Additional resources 334

Dell OpenManage DRAC Tools (RACADM)

An integrated Dell Remote Access Controller (iDRAC) with Lifecycle Controller is embedded in every Dell PowerEdge server. The RACADM command-line utility provides a scriptable interface that enables you to configure the iDRAC either locally or remotely. The utility runs on the management station and the managed system.

The RACADM utility supports the following interfaces:

- Local - Supports running RACADM commands from the managed server's operating system. To run local RACADM commands, install the OpenManage software on the managed server. Only one instance of Local RACADM can be executed on a system at a time.
- SSH or Telnet (also known as Firmware RACADM) - Firmware RACADM is accessible by logging in to iDRAC using SSH or Telnet.
- Remote - Supports running RACADM commands from a remote management station such as a laptop or desktop running Windows or Linux. To run Remote RACADM commands, install the DRAC Tools utility from the OpenManage software on the remote computer.

Note

For information regarding RACADM commands, see the RACADM CLI Guide.

Update the hardware using remote RACADM

You can install and execute the Dell RACADM tool from any management system with access to the iDRAC network. The remote RACADM command set is useful in this situation to mount and execute the Hardware ISO to a large number of VxFlex Ready Node servers. This prevents the need for multiple web browser windows and manual keyboard and mouse clicks for each server.

Note

For additional instructions regarding version updates following SATADOM replacement on 13G servers, see Update the SATADOM firmware (13G servers only) on page 254.

Procedure

1. Configure a laptop or server with access to the iDRAC network.
2. Install the Dell DRAC tools, including RACADM:
   a. In any Internet browser, go to the RACADM Command Line Interface for DRAC page in the Dell Systems Management wiki.
   b. Follow the instructions for downloading and installing Remote RACADM for your operating system.
3. Download the VxFlex Ready Node Hardware Update Bootable ISO from the https://support.emc.com/products/41077 https://support.emc.com/products/42216 and make it accessible on the network share folder.
4. Create a CIFS or NFS network share that includes the Hardware ISO. The share will need to be accessible to the iDRAC network.

   Refer to OS vendor guidelines for instructions on how to create CIFS (windows) or NFS (Linux or XenServer) shares

5. Open a terminal/command prompt with superuser/administrator accessibility. This prompt should allow execution of the RACADM utility.
6. Mount the Hardware ISO to the iDRAC from the remote share, where the following command is all on one line:

```
racadm -r <dracIP> -u root -p <password> remoteimage -c -u
<myuser> -p <mypass> -l //<myip>/SIORN/VxFlex-Ready-Node-
Hardware-Update-for-Dell_14G_2018_May_A00.iso
```

Where:

- *<dracIP>* is the iDRAC IP address
- *<password>* is the password for the server
- *<myuser>* is the NFS/CIFS share username
- *<mypass>* is the NFS/CIFS share password
- *<myip>* is the NFS/CIFS server IP

Note

The default password is calvin for 13G nodes and Scaleio123 for 14G nodes.

7. Enable the iDRAC VirtualCD to boot once using these Remote RACADM commands:

```
racadm -r <dracIP> -u root -p <password> set
iDRAC.ServerBoot.BootOnce Enabled
racadm -r <dracIP> -u root -p <password> set
iDRAC.ServerBoot.FirstBootDevice VCD-DVD
```

8. Verify that the server is in a state in which it can be power-cycled. (The OS should be in maintenance mode, as needed).
9. Power-cycle the VxFlex Ready Node server:

```
racadm -r <dracIP> -u root -p <password> serveraction
graceshutdown
```

10. Repeat steps 6 on page 327 - 9 on page 327 for each VxFlex Ready Node server needing the BIOS and firmware updates and configuration. All servers can be updated in parallel.
11. Wait for the configuration and firmware updates to complete. The server console screen will indicate when the script is complete.

⚠ WARNING

Do not reboot the VxFlex Ready Node server while the update process is being performed!

The iDRAC will be reset several times during the update process. This causes the iDRAC virtual console viewer to close, virtual media to disconnect, and the iDRAC browser window to be unavailable for several minutes during each reset. The hardware update scripts will continue to run from RAM on the server.

The update script will generate a log indicating whether each configuration and firmware flash is successful.

12. (Optional) Check each VxFlex Ready Node server's log for successful competition:
    a. Connect to the iDRAC KVM console screen.
    b. After the update script completes, press Alt+F2 to access a user console, and then Enter to log in.
    c. Check the log contents for errors:

    ```
    less /bundleapplicationlogs/apply_components.log
    ```

    For more information, see Troubleshooting the Hardware ISO on page 331 .

    d. Press q to exit the log viewer, and then Alt+F1 to access the original console screen.
13. Reboot the servers and allow the update and configuration jobs to complete. You can do this remotely using RACADM (see Step 9 on page 327) or locally from the server console screen.
14. For each server, check the iDRAC job queue to verify that the iDRAC jobs created by the bootable ISO have completed successfully:

    ```
    racadm -r <dracIP> -u root -p <password> jobqueue view
    ```

    Any job failures may require running the bootable ISO again, or further troubleshooting.

15. For each server, after the hardware updates are finalized, clean up the iDRAC job queue:

    ```
    racadm -r <dracIP> -u root -p <password> jobqueue delete --all
    ```

16. Continue with OS installation and configuration, VxFlex OS deployment, and other required tasks.

Update the SATADOM firmware (13G servers only)

The Dell factory should provide the SATADOM boot device with the correct firmware already loaded. However, in some cases, it may be necessary to update the firmware loaded on the SATADOM boot device.

For convenience, SATADOM firmware version S130710K is included on the VxFlex Ready Node Hardware ISO.

⚠WARNING

The SATADOM firmware update process permanently erases all data from the SATADOM device. Proceed with caution!

Procedure

1. Connect to the iDRAC Virtual KVM console.

   For instructions, see Open the KVM console - 13G servers on page 318.

2. Boot the VxFlex Ready Node server to the Hardware ISO.

   For instructions, see Update the hardware using the iDRAC virtual console on page 253.

3. After the update script completes, press Alt-F2 to access a user console.
4. Press Enter to login.
5. Execute the SATADOM firmware:

   ```
   /tmp/S130710K_read_log_issue/satadom.sh
   ```

6. Press y to begin the SATADOM update.
7. Press Alt+F1 to access the original console screen, and then press Enter to reboot the server.

Recommended BIOS and firmware settings

This section describes the BIOS, firmware, and configuration settings included in the Hardware ISO.

VxFlex OS ID module

The Hardware ISO runs a script that automatically flashes the ID module, as needed.

This action rebrands a PowerEdge server naming, and is intended to make it compatible with the VxFlex Ready Node offering.

The VxFlex Ready Node ID module should be installed only on VxFlex Ready Node servers. In the event of a system board failure, the Hardware ISO can assist with the reinstallation of the ID module on the replacement system board.

BIOS and firmware

The Hardware ISO runs a script that automatically forces the server to install the necessary firmware updates.

This firmware is consistent with the qualified Dell EMC VxFlex Ready Node Driver and Firmware Matrix, located at https://support.emc.com/products/41077https://support.emc.com/products/42216.

Some of the firmware listed in the table is dependent on the VxFlex Ready Node hardware configuration. The Hardware ISO attempts to apply all firmware updates, but only those updates that are compatible will be installed.

Applying settings using RACADM

The individual firmware files are also available on the Dell EMC Online Support site, and can easily be installed using the following remote RACADM command:

```
racadm -r <dracIP> -u root -p <password> update -f <filename.exe>
```

Where:

- *<dracIP>* is the iDRAC IP address
- *<password>* is the password for the server
- *<filename.exe>* is the name of the Dell Windows update packages

Note

The default password is calvin for 13G nodes and Scaleio123 for 14G nodes.

Configuration settings

The Hardware ISO runs a script that automatically configures the BIOS and iDRAC settings listed in the table below. Some settings are dependent on the server model.

Table 35 Hardware ISO configuration settings

Description	Setting	Value
Server BIOS Boot Sequence	BIOS.BiosBootSettings.BootSeq	HardDisk.List.1-1
Hard Disk Boot Order	BIOS.BiosBootSettings.HddSeq	Disk.SATAEmbedded.J-1,RAID.Integrated.1-1 or Disk.SATAEmbedded.D-1,RAID.Integrated.1-1
Server Boot Mode	BIOS.BiosBootSettings.BootMode	Bios
SRIOV Global Enablement	BIOS.IntegratedDevices.SriovGlobalEnable	Enabled
Memory Performance Tuning	BIOS.MemSettings.SnoopMode	EarlySnoop
System Change Tracking	BIOS.MiscSettings.InSystemCharacterization	Disabled
CPU Virtualization Features	BIOS.ProcSettings.ProcVirtualization	Enabled
CPU Cores	BIOS.ProcSettings.ProcCores	All
CPU X2APIC Mode	BIOS.ProcSettings.ProcX2Apic	Enabled
CPU Turbo Engagement	BIOS.ProcSettings.ControlledTurbo	Disabled
System Power Profile	BIOS.SysProfileSettings.SysProfile	PerfOptimized
OS to iDRAC Pass-through Mode	iDRAC.OS-BMC.PTMode	usb-p2p
OS to iDRAC Enablement	iDRAC.OS-BMC.AdminState	Enabled
iDRAC DHCP Enablement	iDRAC.IPv4.DHCPEnable	Disabled

Table 35 Hardware ISO configuration settings (continued)

Description	Setting	Value
iDRAC Default Credentials	iDRAC.Tuning.DefaultCredentialWarning	Disabled
iDRAC IPMI Enablement	iDRAC.IPMILan.Enable	Enabled
iDRAC Alert Enablement	iDRAC.IPMILan.AlertEnable	Enabled
iDRAC IPv6 Enablement	iDRAC.IPv6.Enable	Enabled

Applying settings using RACADM

The individual settings can also be applied using the remote RACADM command:

```
racadm -r <dracIP> -u root -p <password> set <setting> <value>
```

Where:

- *<dracIP>* is the iDRAC IP address
- *<password>* is the password for the server
- *<setting>* is the BIOS/iDRAC setting name
- *<value>* is the BIOS/iDRAC setting value

Note

When setting the BIOS configuration, include this command:

```
racadm -r <dracIP> -u root -p <password> jobqueue create BIOS.Setup.1-1
<value>
```

Troubleshooting the Hardware ISO

This section describes troubleshooting procedures for problems you may encounter while using the Hardware ISO.

Troubleshoot general iDRAC failures

When problems occur with iDRAC or Lifecycle Controller jobs, you can delete all jobs with a single iDRAC command. All of the completed jobs, plus any orphaned pending jobs, are deleted, and the data manager service on the iDRAC is restarted.

Procedure

1. Clear the iDRAC job queue:

   ```
   racadm -r <dracIP> -u root -p <password> jobqueue delete -i
   JID_CLEARALL_FORCE
   ```

   Where:

   - *<dracIP>* is the iDRAC IP address

- *<password>* is the password for the server

Note

The default password is calvin for 13G nodes and Scaleio123 for 14G nodes.

2. Wait 120 seconds.

   The iDRAC is unable to process any other jobs during this time.

3. Reset the iDRAC

```
racadm -r <dracIP> -u root -p <password> racreset
```

   The iDRAC becomes accessible on the network 3 to 5 minutes after the reset.

iDRAC virtual console issues

Within the iDRAC virtual console window, if the keyboard or some of the keys are not responding, perform the following checks:

- Ensure on your keyboard that the Scroll Lock button is off.
- In the iDRAC virtual console, ensure that the **Keyboard/Mouse Attach State** is set to **Auto-attached**.

For persistent iDRAC virtual console issues, see Troubleshoot general iDRAC failures on page 331.

iDRAC virtual media issues

The following are solutions to problems that may occur when using iDRAC virtual media:

1. If the iDRAC is stuck, one of these actions may assist with recovery:
   - In the server's **Attached Media** screen, disconnect the iDRAC Remote File Share from the iDRAC browser.
   - Disconnect the iDRAC remote image:

```
racadm -r <dracIP> -u root -p <password> remoteimage -d
```

     Where:

     - *<dracIP>* is the iDRAC IP address
     - *<password>* is the password for the server

Note

The default password is calvin for 13G nodes and Scaleio123 for 14G nodes.

2. Issues when booting to iDRAC virtual media may cause CPU machine check errors at POST. When this occurs, clear the iDRAC job queue and reset the iDRAC, as described in Troubleshoot general iDRAC failures on page 331.

DTK A08: BIOS set commands fail

When running DTK A08 on 13G servers with iDRAC FW version 2.50.50.50 or higher, the BIOS set commands may fail.

The DTK is unable to configure the BIOS settings, and the following error message may appear:

```
ERROR: BOOT016: Input source argument value for the boot device is
incorrect or not found among the boot devices on the system.
```

To resolve the issue, configure the BIOS on the server as explained in "Set up the BMC (iDRAC) IP address and BIOS" in this document.

Rebranding operation fails

When updating the Personality Module on 14G servers using the DTK, the operation may fail and the following error message may appear:

```
UEFI0052 Unable to complete the rebranding operation.
```

To resolve the problem, upload the PM file manually:

1. Contact Customer Support for the VxFlex OS PM file.
2. Follow the instructions in the *VxFlex OS Upgrade Guide* for "Single server firmware and BIOS upgrade to a specific version" to upload the file.

Check the logs for error messages

You can view VxFlex Ready Node Hardware Update Bootable ISO logs after the update script completes.

Procedure

1. Press Alt-F2 to access a user console, and then press Enter.
2. Open the log to check the contents for errors:

   ```
   less /bundleapplicationlogs/apply_components.log
   ```

3. You can also view the script for the Hardware ISO, which is useful in helping to identify and troubleshoot log entries:

   ```
   less /opt/dell/toolkit/systems/drm_files/apply_bundles.sh
   ```

Results

The script attempts to configure several boot order commands, regardless of the hardware configuration of the VxFlex Ready Node server. This allows the script to support multiple hardware platform configurations.
Therefore, it is normal to see these error messages within the logs:

```
Apply R730xd bootorder
[Key=BIOS.Setup.1-1#BiosBootSettings]
RAC1017: Successfully modified the object value and the change is in
       pending state.
       To apply modified value, create a configuration job and reboot
       the system. To create the commit and reboot jobs, use "jobqueue
       command. For more information about the "jobqueue" command, see
       help.

ERROR: BOOT016: Input source argument value for the boot device is inc
        not found among the boot devices on the system.

ERROR: BOOT016: Input source argument value for the boot device is inc
        not found among the boot devices on the system.

ERROR: BOOT016: Input source argument value for the boot device is inc
        not found among the boot devices on the system.

[Key=BIOS.Setup.1-1#BiosBootSettings]
RAC1017: Successfully modified the object value and the change is in
       pending state.
:
```

Firmware updates may also display and log the following message:

```
This update is not compatible with your system configuration.
```

```
iDRAC Settings Complete
Proceeding with server BIOS and firmware flash.
Collecting inventory...
Running validation...

This Update Package is not compatible with your system configuration.

Collecting inventory...
Running validation...
```

These firmware are included on the ISO in order to support various VxFlex Ready Node platforms. The log messages do not necessarily indicate a failure.

Additional resources

This section contains information regarding additional resources that may be helpful useful when using the Hardware ISO.

VxFlex OS resources

VxFlex Ready Node deployments have specific guidelines regarding server installation, rack and stack procedures, and power and networking requirements.

The *VxFlex Ready Node Hardware Installation Guide* describes how to install the physical components of a VxFlex Ready Node system. For additional information regarding the VxFlex Ready Node product, documentation, advisories, downloads, and white papers, visit https://support.emc.com/products/41077https://support.emc.com/products/42216.

Dell Lifecycle Controller (LC)

With the launch of the Dell PowerEdge 13th-generation servers in September 2014, Dell enhanced embedded management without the need to install a software-based

agent within the host operating system. 14-generation servers continue this functionality.

At the heart of the PowerEdge servers' embedded management is the iDRAC with Lifecycle Controller (LC) technology. This technology allows users to perform useful tasks such as configuring BIOS and hardware settings, deploying operating systems, updating drivers, changing RAID settings, and saving hardware profiles. Together, they provide a robust set of management functions that can be leveraged throughout the entire server lifecycle.

Supported versions are iDRAC8 (13G servers) or iDRAC9 (14G servers).

For more information, visit the Lifecycle Controller wiki homepage.

Dell OpenManage Deployment Toolkit (DTK)

The Dell OpenManage Deployment Toolkit (DTK) includes a set of utilities, sample scripts, and sample configuration files that you can use to deploy and configure Dell systems.

You can use the DTK to build script-based and RPM-based installation for deploying large number of systems on a pre-operating system environment in a reliable way, without changing their current deployment processes. Using DTK you can install operating systems on Dell systems in BIOS mode (13G) or BIOS or Unified Extensible Firmware Interface (UEFI) mode (14G).

For more information, visit the Dell OpenManagement Deployment Toolkit wiki homepage.

GLOSSARY

A

Active Directory Active Directory (AD) provides directory-based identity-related services. It maintains a directory that is used to centrally store identity information and security principles, and uses them to authenticate and authorize users and devices.

Active Forward Rebuild A copy of stored data is currently being rebuilt on another server, due to planned or unplanned shutdown of a server.

B

Backward Rebuild Data is rebuilt on servers that went offline and became active again. Forward rebuilds can take a long time, and therefore, it can be quicker to restore and update the data on a server which has come back online, than it is to do an entire rebuild on a different server.

Brick A chassis container supplying disk, power and cooling services to a node.

BWC Bandwidth counters.

C

Cache Cache is random access electronic storage used to retain frequently used data for faster access by the channel. Cache is a critical aspect of storage performance. VxFlex OS uses server DRAM for Read RAM Cache (RMcache), as well as SSD/Flash devices (RFcache) for caching reads. VxFlex OS cache uses recently-accessed (LRU) data readily available to manage caching. I/Os read from cache have a lower response time than I/Os serviced by the drives. In addition, cached I/Os reduce the data drive workload, which in many cases is a performance bottleneck in the system.

CacheCade Read and Write caching of storage devices performed by one or more designated SSD devices in the VxRack Node 100 Series system.

Cache Hit Rate The percentage of I/Os from cache.

Cache Skip Data is written directly to storage, bypassing the cache. Reasons for cache skips include: I/Os were too large, the cache device was busy, or I/Os were unaligned. The cache can also be configured to always work in passthrough mode.

Cache Writes Handling Mode The caching write-mode used by the system: passthrough mode (writes to storage only), or cached mode (by default, writes both to cache and to storage).

Cluster Mode VxFlex OS is controlled by a cluster of MDM nodes, minimally consisting of a Master MDM, Slave MDM, and a Tie Breaker node. 5-node clusters consist of one Master MDM, two Slave MDMs, and two Tie Breakers.

D

Degraded Capacity The capacity is available, but is not protected in case of another failure

Device Physical storage device, such as a flash drive, or magnetic disk

DirectPath In VxFlex OS documentation, we use the term DirectPath to refer to the VMware vSphere VMDirectPath I/O feature.

DRL Dirty Region Logging: DRL bits indicate if data is in-writing to a certain location. Once the data is written in both primary and secondary locations, the DRL bit associated with the written location is cleared. These bits can be either stored in DRAM only (memory_only) or also backed up in non-volatile memory (hardened). The former delivers better I/O performance; the latter reduces data movement following a power-cycle giving rise to a faster rebuild.

F

Failed Capacity The capacity is inaccessible due to a failure, and data integrity is at risk

Fault Sets A logical entity that ensures that SDS data is backed up on SDSs that belong to other Fault Sets, thus preventing double-point-of-failure scenarios if rack power outages occur.

Forward Rebuild Data in storage will be rebuilt on another server, due to planned or unplanned shutdown of a server.

I

ID Identifier, a unique sequence of characters that identifies an object in the system. In some CLI commands, an ID can be used to specify a system component.

IP Role The role of the IP address configured for an SDS. Each SDS can have several IP addresses associated with it. Each IP address can serve a different purpose, or role. IP roles include: SDS, SDC, or both SDS and SDC.

L

LDAP The Lightweight Directory Access Protocol (LDAP) is a directory service protocol that runs on a layer above the TCP/IP stack. It provides a mechanism used to connect to, search, and modify Internet directories using Client-Server architecture. In VxFlex OS, LDAP is the protocol used by the MDM to communicate with Active Directory (AD) for authentication purposes.

Lockbox Lockbox is a component of the RSA Common Security Toolkit (CST) which securely stores data (such as passwords) in an encrypted file. A lockbox must be defined for LDAP (secure LDAP), SNMP, and ESRS. For LDAP, lockbox use is optional.

M

Management IPs The IP addresses of the MDMs defined in the system that can be used to access the MDM from VxFlex OS GUI, CLI, and REST.

Management Port The Port number used by the MDM for purposes of communicating with the nodes in the VxFlex OS network.

Manager MDM An MDM that can act as a Master or a Slave in the cluster. Manager MDMs have a unique system ID, and can be given unique names. A manager can be a standby or a member of the cluster.

Master MDM The MDM in the cluster that controls the SDSs and SDCs.

MDM Any server with the MDM package installed on it. An MDM can be given a Manager or a Tie Breaker (default) role, during installation. MDMs have a unique MDM ID, and can be given unique names.

N

Node A node is a server that resides in a chassis in the following form factors:

2U 1 Node—the node occupies the entire space of a 2U-height chassis, and is the equivalent of a standard rack-mounted server.

2U 4 Node—the nodes can be supplied with various processor, RAM, and I/O profiles. Each node occupies a specific slot in the 2U-height chassis and uses 6 disks out of 24 in the chassis. Each node is independent of the others in the same chassis (similar to blade servers).

P

Page Size The page size, typically in KB, used for caching purposes by Read Flash Cache.

Pass-Through Mode Data is passed through to or from storage devices without being cached by Read Flash Cache.

Pending Backward Rebuild A backward rebuild is waiting in a queue, and will be performed when possible, according to rebuild throttling policy.

Primary MDM See Master MDM.

Protected Capacity Capacity that has an accessible copy in the system, in case of failure.

Protection Domain A unique set of SDSs grouped together for reliability and tenancy separation.

R

RAM Read Cache (RMcache) Server RAM that is reserved for caching storage devices in a Storage Pool.

Rebalance When VxFlex OS detects lopsided use of storage capacity, or when new nodes are added, it redistributes data across the nodes, in order to improve performance.

Rebuild When VxFlex OS detects a failure in the network, it creates a new copy of the data from the failed component, in a new location, to ensure data integrity.

Restricted MDM Mode A mode set in which commands can only be performed from an MDM machine.

Restricted SDC Mode Only approved SDCs can access the MDM. When this mode is enabled, volumes can only be added to approved SDCs.

RFcache Read-only caching of storage devices performed on one or more designated SSD devices in a VxFlex OS node.

S

SDBG The SDBG is a technical support troubleshooting tool, used to investigate for "live" systems that retrieves internal information from different VxFlex OS components.

SDC Storage Data Client, a lightweight device driver that exposes VxFlex OS volumes as block devices to the application residing on the same server on which the SDC is installed.

SDS Storage Data Server, which manages the capacity of a single server and acts as a back-end for data access. The SDS is installed on all servers contributing storage devices to VxFlex OS.

Secondary MDM See Slave MDM.

Single Mode A single MDM manages the VxFlex OS network. This mode has no backup protection, and should not be used in production environments.

Slave MDM An MDM in the cluster that is ready to take over the Master MDM role if ever necessary.

Snapshot Capacity The amount of capacity occupied by snapshots of volumes.

Spare Capacity Capacity that is reserved for system use, when recovery from failure is required. This capacity cannot be used for storage purposes.

Spare Percentage Policy This policy determines the amount of capacity that must always be reserved as free space.

Standby MDM An MDM node that is ready to use, with an ID, that has been locked to a specific VxFlex Ready Node system.

Storage Pool A sub-set of physical storage devices in a Protection Domain. Each storage device can only belong to one Storage Pool. User volumes will always use the storage of a single Storage Pool.

T

Thick Capacity Capacity allocated for thick volumes.

Thick Provisioned Volume In virtual storage, thick provisioning is a type of storage allocation in which the amount of storage capacity on a disk is pre-allocated on physical storage at the time the disk is created, meaning that the volume has all its capacity pre-allocated on creation.

Thin Capacity Capacity allocated for thin volumes.

Thin Provisioned Volume Thin provisioning is a method of optimizing the efficiency with which the available space is utilized in storage area networks (SAN). Thin provisioning operates by allocating disk storage space in a flexible manner among multiple users, based on the minimum space required by each user at any given time.

Throttling Throttling controls resource prioritization for rebuild and rebalance processes. Throttling can be controlled per Protection Domain or per Storage Pool (by configuring rebuild and rebalance policies).

Tie Breaker The Tie Breaker (TB) is an MDM that does not have a manager role, whose sole purpose is to help determine which MDM module is the manager that will become the master MDM and take control over the VxFlex Ready Node cluster.

The Tie Breaker ensures that there will always be one Master MDM achieving cluster quorum. In a 3-node cluster, there is one TB; in a 5-node cluster, there are two TBs.

U

Unavailable Capacity Capacity that is not being used, but is also unavailable (due to server outage).

Unused Capacity Capacity that is not currently being used for any purpose in the system.

V

Volume A general term referring to a storage device. In VxFlex OS, a volume consists of multiple blocks spread evenly on Storage Pool devices.

W

Widget The full screen view can be minimized into a widget, which is a small window that floats on your screen, over other applications. Property sheets can also be minimized into widgets.

Write Misses Write requests that were not found in cache