Dell EMC VxFlex Ready Node

Version 2.x

AMS Deployment Guide

302-003-032

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 7

Tables 9

Preface 11

Part 1 Introduction 13

Chapter 1 Introduction to VxFlex Ready Node 15
What is VxFlex OS? 16
System requirements 17
Server requirements 17
Other requirements 18
Boot drive 19
Switch requirements 19
IP addresses 19
New and changed features 19
New features for VxFlex OS 20
New features for VxFlex Ready Node 20
Product limits 20

Chapter 2 Architecture 23
VxFlex OS architecture overview 24
System 24
Hardware 24
Software 24
The MDM cluster 25
Storage definitions 27
Protection Domains 27
Storage Pools 27
Fault Sets 29
Naming 30
Protection and load balancing 30
Rebuild 30
Rebuild throttling 30
Rebalance 31
Rebalance throttling 32
Checksum protection 32
ESX vStorage APIs for Array Integration (VAAI) 33
Caching 34
Networking 36
IP Addresses 36
VLAN support 38
Networking and connectivity - 13G servers 39
Networking and connectivity - 14G servers 46
Virtual IP Address 56
Monitoring of SDC and SDS connections 57

Snapshots....57
V-Trees.... 58
Other functions.... 59
Implementing VxFlex OS.... 60
Implementing VxFlex OS in an ESXi-based system.... 60
Implementing VxFlex OS over a Linux bare metal system.... 62
SAN virtualization layer.... 63
Compute cluster configuration and workload deployment.... 64
Maintenance....65
Maintaining the physical layer....65
Instant maintenance mode....66
Maintaining the virtualization layer.... 66
Management tools....67
Firmware patching....67

Part 2 Deploying VxFlex Ready Node 71

Chapter 3 Deploying VxFlex Ready Node 73
VxFlex Ready Node deployment overview.... 74
Deployment setup....75
Install the VxFlex OS GUI....75
Install the VxFlex Ready Node AMS....75
Enable the zero padding parameter.... 76
Log in to the VxFlex OS GUI.... 77
Perform initial AMS setup....77
Password and IP addresses.... 80
VxFlex Ready Node deployment wizard overview....82
Deploy or extend a VxFlex Ready Node system....85
Configuring the number of devices in a Storage Pool.... 88
Extend a system with nodes of a different operating system.... 89
Extend a VMware-based system with non-ESXi nodes....89
Extend a Linux-based system with non-Linux nodes....90

Chapter 4 Post-Deployment Activities 93
Post-deployment checklist.... 94
Create and map volumes.... 94
Creating and mapping volumes using the VxFlex OS GUI....95
Creating and mapping volumes using the CLI.... 96
Update the system component performance profiles.... 96
Configuring ESRS....97
Configure native users....98
Configure LDAP users.... 99
Configuring SNMP using the VxFlex OS GUI.... 100
Setting credentials for existing Monitor users....100
Back up SSH public keys using the VxFlex OS GUI....101
Configure the timing of the alert on certificate expiration....102
Sign and update AMS security certificates.... 102
Activating the ESX license on a node.... 103
Adding an external SDC to an existing VxFlex OS system.... 103
Install the SDC on an ESX server and connect it to VxFlex OS.....103
Installing SDC on a Linux server and connecting it to VxFlex OS..104
Installing SDC on a Windows server and connecting it to VxFlex OS 105
Uninstall VxFlex Ready Node components.... 105

Uninstalling the VxFlex Ready Node AMS from a Windows server.... 106
Uninstalling the VxFlex Ready Node AMS from a Linux server..... 106
Uninstall the VxFlex OS GUI.....106
Prevent certificate errors when ESX hosts are added using hostnames....107

Chapter 5 Licensing 109
Licensing overview..... 110
Activating entitlements and installing a license file..... 111
Activate an entitlement and download the license file..... 111
Install the license - using the VxFlex OS GUI.....115
Install the license - SCLI..... 115
License file example.....116
Error messages.....116

Part 3 Reference 119

Chapter 6 Common Tasks 121
Associating VxFlex OS volumes with physical disks.....122
Volume information - Linux..... 122
Port usage and changing default ports.....123
Upgrading Java..... 124
Move ESXi hosts to a new vCenter and reconfigure AMS.....124
Best practices for using VMware vMotion..... 125
Disable VMware Admission Control..... 126
Update the hardware using the iDRAC virtual console.....127
Update the SATADOM firmware (13G servers only).....128
AMS management takeover of non-managed nodes..... 129
AMS management takeover compatibility..... 129
AMS management takeover of non-managed nodes..... 130

Chapter 7 Server factory reset and power operations 137
Perform a factory reset.....138
Set up the BMC (iDRAC) IP address and BIOS..... 138
Verify the status of the system hardware, storage controller, and disks - 13G servers..... 140
Verify the status of the system hardware, storage controller, and disks - 14G servers.....141
Open the KVM console - 13G servers..... 142
Open the KVM console - 14G servers..... 143
Map the ISO file on a VxFlex Ready Node server.....144
Map the ISO file on a VxFlex Ready Node 14G server.....145
KVM activities.....146
Additional tasks using the KVM console - 13G server..... 146
Additional tasks using the KVM console for 14G servers.....147

Glossary 149

FIGURES

1 5-node MDM cluster....27
2 Protection Domains and Storage Pools....28
3 Protection Domains, Storage Pools, and Fault Sets....29
4 Single-node R730xd (2U1N) server (storage only) typical connection diagram....40
5 Single-node R730xd (2U1N) server (converged) typical connection diagram....40
6 Single-node R630 (1U1N) server typical connection diagram....41
7 Correlating physical slot and Ethernet network names....46
8 R640 and R740xd PCI slots, integrated NICs and BMC port locations....48
9 Correlating physical slot and Ethernet network names....56
10 Snapshot operations....58
11 V-Tree diagram....59
12 VxFlex OS implementation on ESX....61
13 System Settings window, Technical Users pane, Monitor User....101
14 Licensing LAC email....111
15 License file example....116
16 Network ports....126

TABLES

1 Server requirements....17
2 Product limits....20
3 MDM cluster modes....26
4 Caching modes....34
5 Caching support matrix....36
6 VxFlex OS management IP network....36
7 VxFlex OS data IP network (continued)....37
8 R730xd single-node server configuration....41
9 R630 single-node server configuration....42
10 R630 single-node server configuration....43
11 R730xd single-node server configuration....43
12 R730xd single-node server configuration (Storage Only)....44
13 R740xd server configuration for 10/25 GbE, 32xSAS, SFP+/SFP28....49
14 R740xd server configuration for 10/25GbE, 12xNVMe, SFP+/SFP28....50
15 R740xd server configuration for 25GbE, 3xGPU (M60)/ 2xGPU (M10), SFP28....51
16 R740xd server configuration for 10GbE, 3xGPU (M60)/ 2xGPU (M10), SFP+....51
17 R640 server configuration for 10/25GbE, 1CPU, SFP+/SFP28....52
18 R640 server configuration for 10/25GbE, 2CPU, SFP+/SFP28....52
19 R640 single-node server configuration....53
20 R740xd single-node server configuration for 10/25GbE, 2CPU, SFP+/SFP28....54
21 13G PowerEdge Servers....67
22 14G PowerEdge Servers....68
23 Password and IP address table....80
24 Caching options....83
25 eLicensing terminology(continued)....110
26 Licensing error messages....116
27 Default ports....123

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 VxFlex Ready Node benefits and architecture. 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? 16
- System requirements 17
- New and changed features 19
- Product limits 20

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 (AMS) server	• Physical requirements: ▪ Processor - 2 vCPUs ▪ Memory- 3 GB ▪ Storage - 80 GB of free storage capacity • Operating systems (all 64-bit): ▪ AMS installed on a VM hosted on ESXi servers must use the vmxnet3 vmnic. Use of the e1000 vmnic is not supported, and causes network failures. ▪ 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. ▪ 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

Table 1 Server requirements (continued)

Component	Requirement
	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
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)
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 36 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 24
- System 24
- The MDM cluster 25
- Storage definitions 27
- Protection and load balancing 30
- ESX vStorage APIs for Array Integration (VAAI) 33
- Caching 34
- Networking 36
- Virtual IP Address 56
- Monitoring of SDC and SDS connections 57
- Snapshots 57
- Other functions 59
- Implementing VxFlex OS 60
- SAN virtualization layer 63
- Compute cluster configuration and workload deployment 64
- Maintenance 65
- Management tools 67
- Firmware patching 67

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

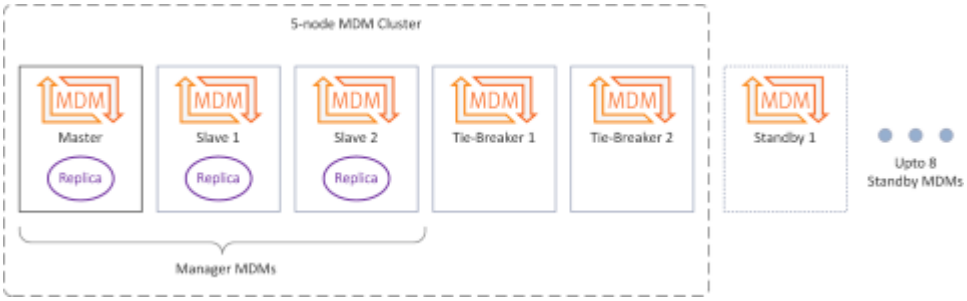

All members of the MDM cluster have the same MDM package installed on them.

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

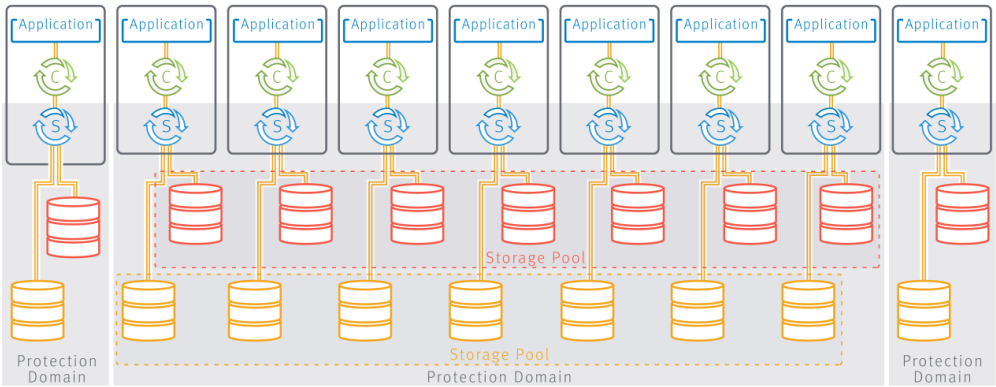

All devices in a Storage Pool must be of the same device media type.

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 29 illustrates the same configuration as Figure 2, with the addition of Fault Sets.

Figure 3 Protection Domains, Storage Pools, and Fault Sets

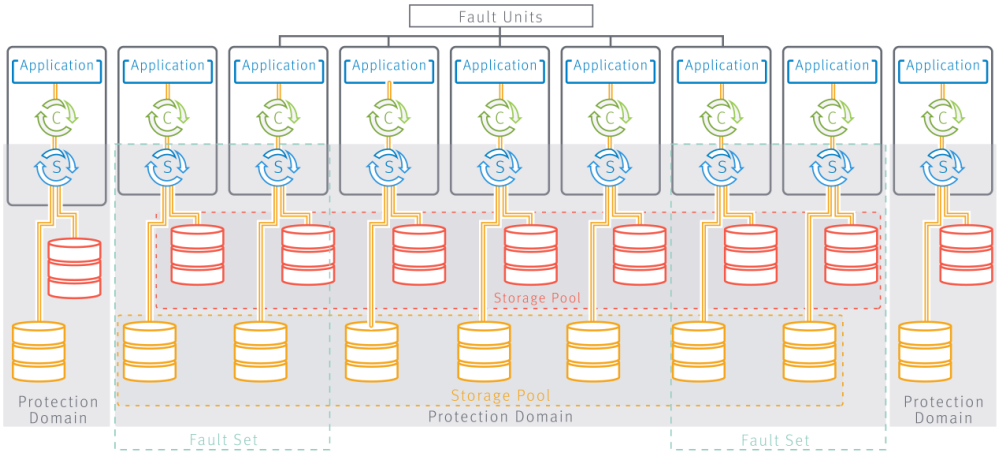

To use Fault Sets, you must work in the following order:

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.

Note

AMS running on a VM hosted on ESXi servers must use the vmxnet3 vmnic. Use of the e1000 vmnic is not supported, and causes network failures.

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.

Table 6 VxFlex OS management IP network (continued)

Item	Description	Comments
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
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	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
	3. MDM_Cluster_Virtual_IP_DATA1 = the Virtual IP addresses of the MDM cluster in the "Data1" network	
	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

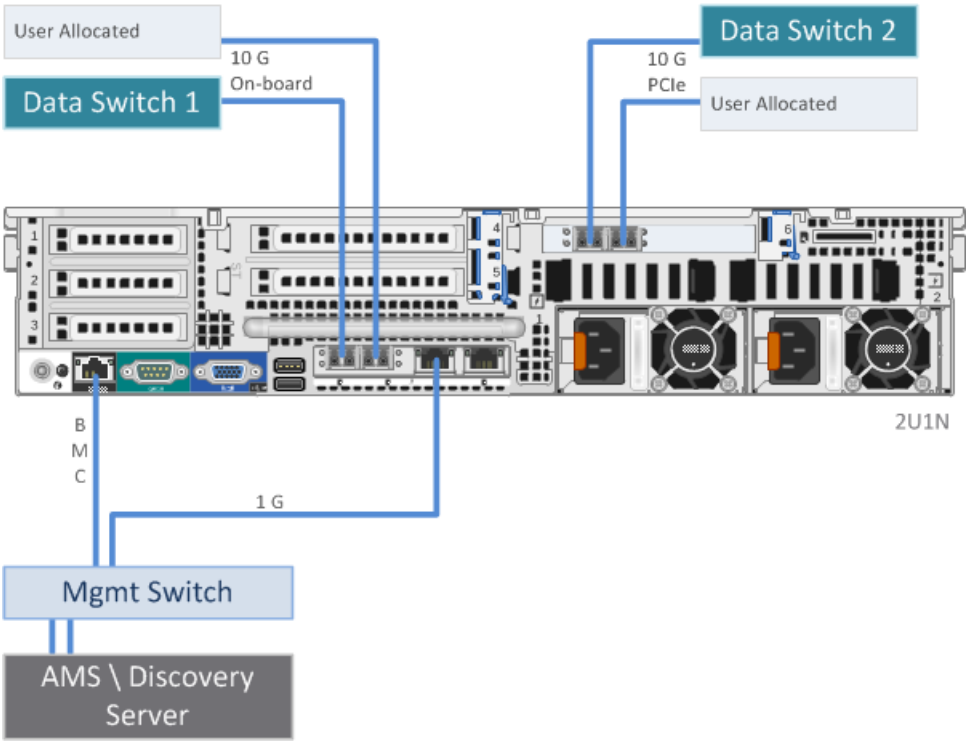

Figure 5 Single-node R730xd (2U1N) server (converged) typical connection diagram

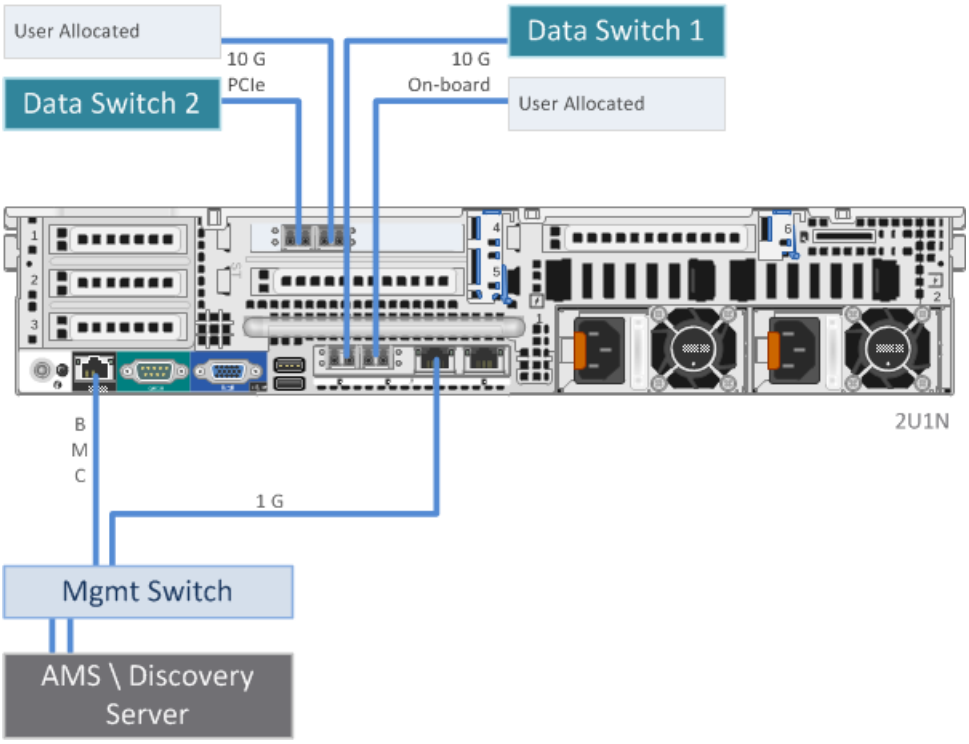

Figure 6 Single-node R630 (1U1N) server typical connection diagram

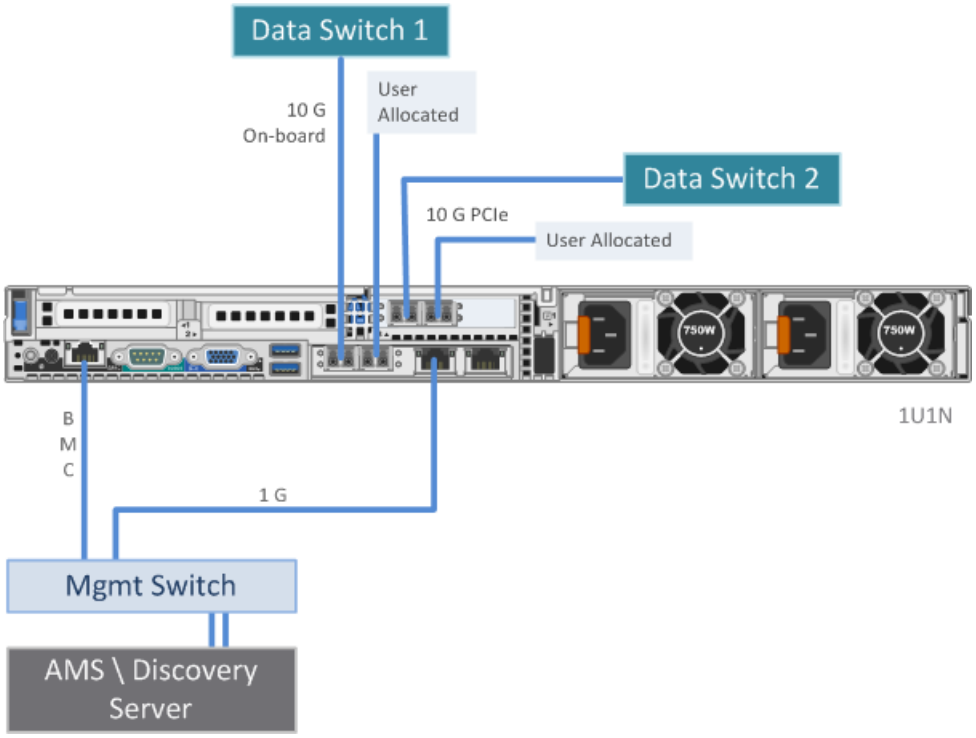

Server node port designations - ESXi 13G servers

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 40 and Figure 5 on page 40.

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 46, 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 42.
   - 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 44 output to the corresponding line in Step 1.b on page 45 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 44 and 1.b on page 45:

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

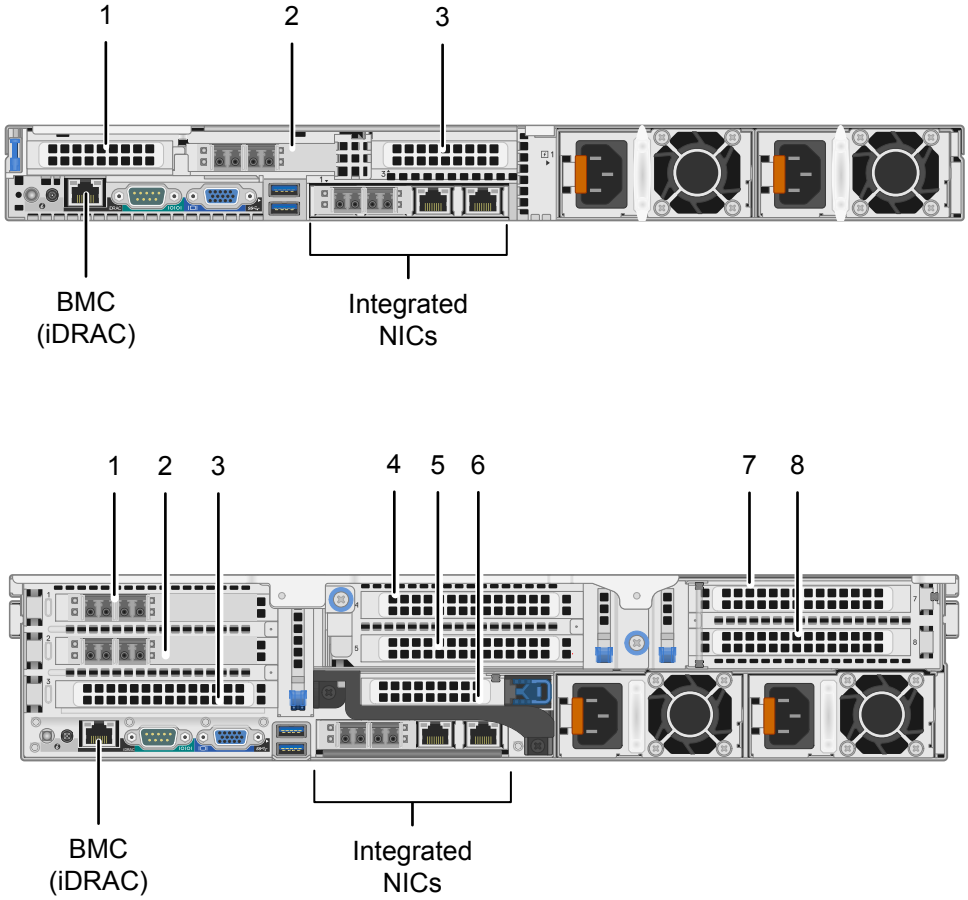

Data cables are connected to two switches for high availability, via the nodes' 10 GbE, 25 GbE, or 100 GbE ports. Management interfaces are connected to a switch on a separate management network using onboard and BMC (iDRAC) ports.

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 Management Network	Not in Use
vSwitch	N/A	sio_vSwitch1	User allocated	User allocated	sio_vSwitch2	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 Management 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_vSwitch2	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 Management 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_vSwitch1	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 54 output to the corresponding line in step 1.b on page 55 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 54 and 1.b on page 55:

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

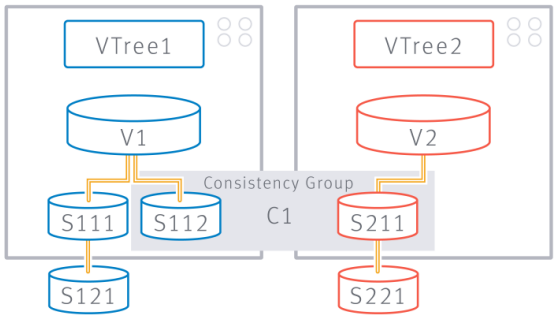

The structure related to all the snapshots resulting from one volume is referred to as a V-Tree (short for volume tree). When taking a snapshot in the system, you can specify more than one volume. All snapshots taken together form a consistency group. They are consistent in the sense that they were all 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 58, 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

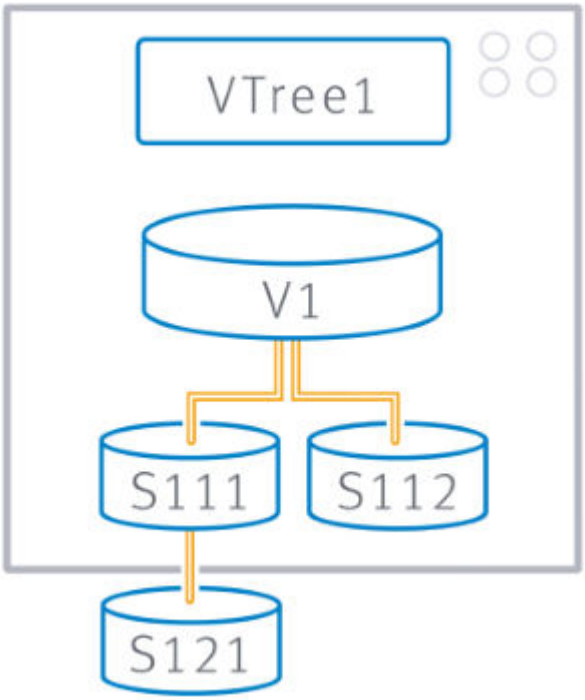

Other functions

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

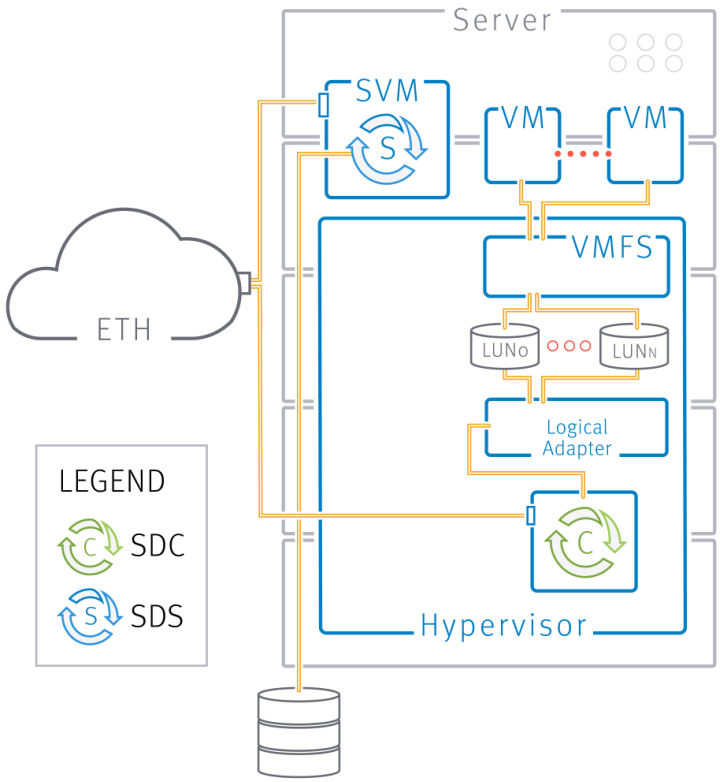

The LUNs in the previous figure can be formatted with VMFS, and then exposed using the ESXi host to the virtual machine, or can be used as RDM devices. When the LUNs are used as RDM devices, the VMFS layer is omitted.

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 29.

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

Deploying VxFlex Ready Node

This part describes how to deploy VxFlex Ready Node in your environment.

Chapters include:

Chapter 3, "Deploying VxFlex Ready Node"

Chapter 4, "Post-Deployment Activities"

Chapter 5, "Licensing"

CHAPTER 3

Deploying VxFlex Ready Node

This section contains procedures on how to install, deploy and perform initial configuration of VxFlex Ready Node software components on VxFlex Ready Node servers.

- VxFlex Ready Node deployment overview 74
- Deployment setup 75
- VxFlex Ready Node deployment wizard overview 82
- Deploy or extend a VxFlex Ready Node system 85
- Extend a system with nodes of a different operating system 89

VxFlex Ready Node deployment overview

Deploying VxFlex Ready Node is performed using the deployment wizard, which is accessed from the VxFlex OS GUI. The VxFlex OS GUI is the client interface to the AMS (Automated Management Services). The wizard discovers nodes in the VxFlex Ready Node system and deploys all necessary software components on the selected nodes.

To deploy VxFlex Ready Node, you must first install the VxFlex OS GUI and AMS. After this initial one-time setup, the deployment wizard guides you through the deployment process.

To deploy and configure nodes the following prerequisites are needed:

- Nodes must be in factory-installed state, with the operating system installed, for both new and replacement deployment. If the node powers on with the operating system, it is in the proper state. If not, follow the instructions in Server factory reset and power operations on page 137 before starting deployment.

  Note

  To avoid deployment conflicts, do not manually add ESXi factory-installed-state nodes to a vCenter environment until deployment is complete; if the vCenter is configured, the deployment process will add them as part of the process.

- Nodes that failed during install, must be reverted to the factory-installed state prior to redeployment. For more information, see Perform a factory reset on page 138.
- Nodes must not be managed by a vCenter.
- During the initial configuration of node IP addresses, the AMS must be on the same physical, layer-two network of the VxFlex Ready Node management network (using the same VLAN). When the AMS cannot be on the same network, you must install a secondary AMS as a proxy discovery server that will be on the node network. The proxy server communicates with the nodes, as well as with the main AMS. The choice of functioning as a proxy server is made during the AMS installation. If the AMS and the nodes are on different subnets, ensure that there is an AMS proxy discovery server installed on the node network. You must also use the VxFlex OS GUI **Preferences** > **System Settings** > **Discovery** option to enable communication with the proxy server.
- The hostname of the AMS server cannot contain an underscore character.
- Ensure that network and other cables are set up according to the description in Networking on page 36.
- Verify that all servers to be added are powered on.
- Zero padding is an optional configuration for Storage Pools. You are required to enable zero padding for specific use cases such as different scenarios for replication. For more information, see Enable the zero padding parameter on page 76.
- The virtual machines on which the active MDM Master or Slave runs act as an NTP server for each SVM/ESX. The SVMs and ESXs connect to the NTP using the MDM cluster's virtual IP address. The MDMs can be configured, during or after deployment, to connect to an external NTP server provided by the user, and the AMS must be configured to connect to the same server. In addition, if you want to use either SSL or LDAP, you must configure an external NTP server. Each

time the external NTP server configuration changes, VxFlex OS will update all of the MDMs with the new NTP IP addresses and initiate NTP synchronization on all of the nodes.

Deployment setup

Before deploying, you are prompted to complete an initial one-time setup that is launched when you log into the VxFlex OS GUI for the first time.

After entering the IP address or hostname of the AMS, you are prompted to approve the security certificate between the VxFlex OS GUI and the AMS. After approving the certificate, you enter the password and IP address ranges. These values serve as pools that form the basis for password and IP address assignment for initial deployment, as well as for adding nodes to the system in the future. You can use Password and IP addresses on page 80 to help you organize the required input.

Before initiating deployment:

- Ensure that you have prepared the passwords and IP address pools described in Password and IP addresses on page 80.
- VLAN tagging is supported to separate the networks on different VLANs. To use VLAN tagging, ensure that the switch in the interface is configured with trunk mode.
- IPv6 is not supported in VMware converged solutions. Contact EMC support for more information.

Install the VxFlex OS GUI

You can install the VxFlex OS GUI.

Before you begin

- Ensure that the workstation satisfies the requirements described in the "System Requirements" section of the documentation.
- Get the installation file either from the product ISO or the Support Site.

You must use the VxFlex OS GUI to deploy a VxFlex Ready Node system with AMS management.

Procedure

1. Install the VxFlex OS GUI:
   - Windows:

```
EMC-ScaleIO-gui-2.6-<build>.X.msi
```

After you finish

To log in to the VxFlex OS GUI, see "Log in to the VxFlex OS GUI."

Install the VxFlex Ready Node AMS

Install the Automated Management Services (AMS). The AMS manages one VxFlex OS system.

Before you begin

- The AMS must be installed on a server with access to the VxFlex Ready Node nodes, but not on a VxFlex Ready Node server.

- Ensure that your AMS server meets the system requirements described in System requirements on page 17.
- Get the installation file either from the product ISO or the Support Site.
- AMS running on a VM hosted on ESXi servers must use the vmxnet3 vmnic. Use of the e1000 vmnic is not supported, and causes network failures.

You can install the AMS on a Windows or Linux workstation or server. During the installation, you can choose to install as the AMS normally, or as a proxy discovery server.

Procedure

1. Run the following command:
   - 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.
2. When using Linux AMS, restart the AMS server.

Enable the zero padding parameter

You can enable the zero padding parameter, which is required for certain replication needs.

By default, the zero padding parameter is disabled in the system. To enable it, you must add the `global.zero_padding=true` parameter at the end of the file.

Note

Ensure that zero padding, when required, is enabled before the start of the deployment process.

Procedure

1. Using a text editor, open the `gatewayUser.properties` file, located in the following directory on the AMS server:
   - Windows 64-bit: `C:\Program Files\EMC\ScaleIO\AMS\webapps\ROOT\WEB-INF\classes\`
   - Linux: `/opt/emc/scaleio/ams/webapps/ROOT/WEB-INF/classes`
2. Add the following parameter at the end of the file:

   ```
   global.zero_padding=true
   ```
3. Save and close the file.
4. Restart the VxFlex OS AMS service:

- Windows: Restart the EMC ScaleIO AMS service
- Linux: Type the following command:

```
service scaleio-ams restart
```

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)

Note

The first time you use the VxFlex OS GUI client, you are prompted for a password. This first pass is a temporary password. Use any string (like "aa"), but not the password that you want to be your permanent password.

Procedure

1. Open the VxFlex OS GUI by selecting **Start** > **All Programs** > **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 AMS username and a password, then click **Login**.

   If this is the first time using this VxFlex OS GUI client, you are prompted to change the temporary password. Type and confirm a new password, then click **Login**.

After you finish

If this is a first-time deployment, the AMS Initial Setup Wizard begins. If you have already populated the IP address pools, skip to Deploy or extend a VxFlex Ready Node system on page 85.

Perform initial AMS setup

Perform a one-time AMS set up before using it to deploy a system.

When you first set up the AMS, an introductory screen appears that describes the process. Review the displayed information, then click **Next** to begin the setup.

Procedure

1. 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 as well as employ relevant validation rules. You can add nodes of a different operating system type after completing the initial deployment.

Note

The default for Linux deployment is a 2-layer architecture. After deployment, manually install the external SDC.

This documentation describes steps for all operating systems, even though some steps (such as SVM configuration) are not relevant for Linux deployment.

c. If you installed a different AMS as a Discovery server (proxy), type its accessible IP address to connect the main AMS to the remote discovery server.

Only 1 discovery server is supported.

d. Click **Next**.

2. In the **Security** screen, define new passwords.

Review minimum password criteria in the Password and IP addresses on page 80 table and keep a record of the passwords for later reference.

a. Type a new host operating system password.

b. If requested, type a new storage Virtual Machine (SVM) password.

c. Type a new BMC (iDRAC) password.

d. Type a new AMS recovery key.

This password is required for running the AMS recovery process.

e. Click **Next**.

3. In the **Management Networks** screen, configure IP address ranges for the management networks.

The range of IP addresses must be at least equal to the number of servers being installed. The management network is used by the AMS to communicate with and control the nodes. No IO pass through this network. The management networks can be on different subnets, as long as the AMS can communicate with them, either by local IP addresses or with a router. If a router is used, the discovery server must be on the same level-2 physical network as the VxFlex OS nodes to enable 2-layer broadcasts.

a. Select an IP address version (IPv4 or IPv6) in **IP address version**.

b. Type IP address ranges for each management network, using the options described below.

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 .

4. In the **Data Networks** screen, configure IP address ranges for each data network.

   Data networks are 10/25/100 GbE networks that are used for inter-node communication and for passing IO.

   a. Select an 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, select that option (for Linux only), and configure the additional NICs.

      You can define two or four NICs.

   c. Type an IP address for the cluster virtual IP.

      The IP address is derived from the host data network.

   d. Type IP address ranges for the data network, using the same options as those described for the management networks.

   e. Type ranges for both the ESX host as well as the storage VM (SVM) that resides on each host.

   f. Click **Next**.

5. (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.

6. (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, and the **Add Nodes** option is displayed in the GVxFlex OS GUIUI.

After you finish

If you performed this procedure as part of an AMS recovery, return to the "Recovering an AMS server" section of the *VxFlex Ready Node Troubleshooting Guide*.

Password and IP addresses

Before deploying, you should prepare the passwords and IP address pools.

The following table describes the inputs needed during the deployment process. Not all fields may be relevant for your operating system environment.

Table 23 Password and IP address table

Field	Default setting	Description/Comments
AMS admin password		The password to be assigned for the default admin user. The password must meet the following criteria: • Between 6 and 31, ASCII-printable characters • No blank spaces • Include at least 3 of the following groups: [a-z], [A-Z], [0-9], special chars (!@#$ ...)
IP address of the auto discovery proxy server		The IP address for the AMS server to serve as an auto discovery server
Host root password	Scaleio123	The password to be assigned for the root user on all hosts in the system. The password 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
SVM root password	admin	The password to be assigned for the root user on all storage VMs in the system.
BMC (iDRAC) root password	password	The password to be assigned for the root user on the BMC (iDRAC) console. The password must meet same criteria as the host root password.
AMS recovery password		The password to be used if an AMS host needs to be recovered in the future. The password must meet same criteria as the host root password.

Table 23 Password and IP address table (continued)

Field	Default setting	Description/Comments
Host management IP pool (with subnet mask and default gateway)		The IP addresses to be assigned for management IPs on the hosts. It is strongly recommended to reserve, and enter a range of IP addresses that is enough to satisfy further extension.
SVM management IP pool (with subnet mask and default gateway)		The IP addresses to be assigned for management IPs on the storage VMs.
Virtual IP for data network A		An IP address assigned to the MDM cluster to enhance communication with the SDS and SDC in the system. The virtual IP address moves with the Primary MDM.
Host data A IP pool (with subnet mask)		The IP addresses to be assigned for data IPs on the hosts, on data network A.
SVM data A IP pool (with subnet mask)		The IP addresses to be assigned for data IPs on the SVMs on data network A.
Virtual IP for data network B		Same as data network A.
Host data B IP pool (with subnet mask)		Same as data network A.
SVM data B IP pool (with subnet mask)		Same as data network A.
Host data C IP pool (with subnet mask) - optional, for Linux only		Same as data network A.
SVM data C IP pool (with subnet mask) - optional, for Linux only		Same as data network A.
Host data D IP pool (with subnet mask) - optional, for Linux only		Same as data network A.
SVM data D IP pool (with subnet mask) - optional, for Linux only		Same as data network A.
BMC (iDRAC) IP pool (with subnet mask and default GW)		
vCenter IP address		The IP address of the vCenter that manages the VxFlex OS nodes on the management network.
vCenter username		A vCenter username with administrative rights. Any valid name can be used, including those with special characters (for example: XXXXX@vSphere.local).

Table 23 Password and IP address table (continued)

Field	Default setting	Description/Comments
vCenter administrator password		The password of the vCenter that manages the VxFlex OS nodes.

VxFlex Ready Node deployment wizard overview

The deployment wizard discovers nodes in the system and deploys all necessary software components on the selected nodes.

Use the deployment wizard, accessed from the VxFlex OS GUI, to perform initial deployment and to add nodes to ("extend") an existing system. You can deploy nodes in the following modes:

- Simple:
  The AMS deploys all available nodes in the system using IP addresses in the IP pool. When you start the simple deployment, all steps are completed with their default values. Changes can be made after deployment completes.

  Note

  In the Linux environment, simple deployment assumes a 2-layer architecture, where external SDCs are installed manually after deployment.

- Advanced:
  The AMS presents a system using IP addresses in the IP pool. You can edit all the node settings before deployment, as well as configure the following:
  - Enable dynamic allocation of unused ports for data networks (1 GbE ports are not supported).
  - Define MDM and Tie-Breaker roles for the nodes.
  - In the Linux environment, determine whether SDCs are deployed on the same servers (converged) or on external servers (installed manually after deployment).
  - Configure the disk role (storage or caching) and what kind of caching (RFcache or DAS Cache).
  - Create and assign nodes to Protection Domains and Fault Sets.
  - Create and assign devices to Storage Pools.

After completing the wizard, several post-deployment actions are suggested.

Note

Deployment must be performed with AMS secured connection enabled (default).

Deployment stages

The deployment wizard follows these stages:

1. Start the deployment wizard.
   When you start the wizard, it discovers those nodes that are configured to listen for an AMS discovery server layer-2 detection broadcast. You can select either a simple or advanced deployment. The simple deployment configures all steps

according to default values. The advanced deployment enables you to review and edit the configuration.

2. Select nodes.
   Select the nodes to be configured, minimally three nodes (two MDM and one Tie-Breaker MDM). The selected nodes are locked to your system and can no longer be discovered by other processes.

   You can filter to locate a node by serial number or service tag.

3. Set management IP addresses.
   Set the management IP addresses to assign to the selected nodes. The IP addresses are assigned from the IP address pool created in the setup stage.

4. Set data IP addresses and assign NICs.
   Set the data IP addresses to assign to the selected nodes and the data ports (NICs) to assign to the data networks. The IP addresses are assigned from the IP pool created in the setup stage.

   For every server, there must be at least two data NICs assigned. In RHEL servers, if there are more data networks, assign one NIC per network. The deployment wizard will attempt to apply a default solution, which you can accept, or you can assign other values.

   Once a single node of a type is configured, you can duplicate its settings to other nodes of the same type. For example, after configuring an R640 node, you can apply its settings to all other R64 nodes, or an R740xd-32 disk chassis can be duplicated to the same.

5. Deploy VxFlex OS components.
   Assign VxFlex OS cluster roles and deploy VxFlex OS software on the nodes.

   The MDM/TB components are installed, according to your assignment. SDS and LIA are installed on every node. On ESXi nodes, SDC is deployed on each node; in RHEL, SDC is not (by default).

6. Add SDS nodes to system.
   Create Protection Domains and Fault Sets, and assign SDS nodes (without their devices) to these objects. By default, all nodes are added to the default Protection Domain.

7. Configure caching.
   Configure the usage of each node's SSD and NVMe SSD devices, either for storage or to provide caching service. The settings are applied consistently for the device. For example, if a device is configured for caching, it will be the same way on all nodes to which that device is connected.

   The caching policy for HDD devices will be set to match what is configured for the SSDs in this step.

   You can make the following caching assignments:

   Table 24 Caching options

Device type	Caching option	Description
SSD	Use for storage	Use for fast storage.
	Use for cache (DAS Cache)	Use to provide DAS Cache caching for HDD devices on the node.

Table 24 Caching options (continued)

Device type	Caching option	Description
		Note A maximum of two SSD devices per node can be configured as DAS Cache, for 13G hybrid nodes only. 13G hybrid nodes were End of Sale as of December 31, 2017; this option is only available in server/component FRU scenarios where the DAS Cache entitlement was part of the original order before the End of Sale date.
	Use for cache (RFcache)	Use to provide RFcache caching for HDD devices on the node.
NVMe SSD	Use for storage	Use for fast storage.
	Use for cache (RFcache)	Use to provide RFcache caching for HDD devices on the node.

If an option is not available for a device, it is not supported by that node's configuration.

8. Add devices to SDS nodes.
   Assign installed devices to Storage Pools (SP). SPs allow the generation of different storage tiers for different forms of storage media in the VxFlex OS system. An SP that contains one form of device cannot be used for other kinds of device.

   Storage Pool creation and device assignment is automatically performed according to default rules. You can use this step to change those assignments.

   You can manually set the maximum number of devices allowed in a Storage Pool. For more information, see Configuring the number of devices in a Storage Pool on page 88.

   The following table displays the default device assignment rules:

Drives per node	Storage Pool 1	Storage Pool 2	Storage Pool 3
<=10	all drives		
11	5	6	
12	6	6	
13	6	7	
14	7	7	
15	7	8	
16	8	8	
17	8	9	
18	9	9	
19	9	10	

Drives per node	Storage Pool 1	Storage Pool 2	Storage Pool 3
20	10	10	
21	7	7	7
22	7	7	8
23	7	8	8
24	8	8	8
25	8	8	9
26	8	9	9
27	9	9	9
28	9	9	10
29	9	10	10
30	10	10	10
31	7	8	8
32	8	8	8

Deploy or extend a VxFlex Ready Node system

Run the deployment wizard to deploy a new system or to add nodes of the same operating system type to an existing system.

Before you begin

Before beginning, ensure the following:

- You have authentication credentials to use the VxFlex OS GUI to access the AMS.
- The AMS is in secure communication mode.

To add nodes of a different operating system type, see Extend a system with nodes of a different operating system on page 89.

Procedure

1. From the VxFlex OS GUI, start the deployment:
   a. Start the deployment wizard, in one of the following ways:
      - If you have just completed the initial AMS setup, click the **Add Nodes** icon.
      - From the VxFlex OS GUI **System Settings** drop-down, click **Add Nodes**.

      The discovery process begins. When it completes, it shows how many nodes are found in the chassis.

   b. Select which deployment mode to use:
      - Simple: Click **Deploy**. The deployment proceeds with the default setup and populates the IP address fields from the IP pool. After deployment completes, skip to the end of this section.
      - Advanced: Click **Advanced**.

   The selected deployment method begins.

2. Select nodes to deploy:
   a. In the **Discover Nodes** screen, select which nodes to add to this VxFlex Ready Node system.

      You can filter by chassis ID or by server type.

   b. You can assign a node's name and a location in the **Name** and **Location** columns.

   c. Click **Select and Proceed**.

   The deployment begins. Progress is displayed in the **Status** column.
3. Set management IP addresses:
   a. In the **Set Management IPs** screen, review the IP addresses that are automatically assigned to the selected nodes. To change the assigned IP addresses, edit the fields.

   b. Click **Set IP Addresses**.

   IP address assignment begins. Progress is displayed in the **Status** column.
4. Set data IP addresses and assign NICs:
   a. In the **Set Data IPs** screen, review the IP addresses that are automatically assigned to the selected nodes. To change the assigned IP addresses, edit the fields.

   b. Review and assign data NICs:

      The deployment wizard will attempt to apply a default solution, which you can accept, or you can assign other values.

      For every server, there must be at least two data NICs assigned. In RHEL servers, if there are more data networks, assign one NIC per network.

      To assign NICs to a data network:

      a. In the **Nodes** pane, select a node.

      b. In the **Data Ports** pane, select an available NIC to assign to the node.

      c. From the **Assign to** drop-down, select to which data network to assign the NIC.

      d. To replicate the NIC assignment to all nodes of the same type, click **Apply to all nodes**.

   c. Click **Set IP Addresses**.

   IP address and NIC assignment begins. Progress is displayed in the **Status** column.
5. Deploy VxFlex OS components:
   a. In the **Deploy** screen, review the nodes and their assigned roles.

   b. To change the node's Cluster role, click the cluster role to open the **Assign Cluster Role** dialog box.

      - Click **Clear** to clear the node's current role.
      - Click **Assign** to assign an available role to the selected node.

   c. On RHEL nodes, where the default is not to install an SDC, you can install SDC by selecting **Install SDC**.

d. Click **Deploy**.

The deploy operation performs node operations in parallel, and can take several minutes to complete. Nodes will be reset during the process. You can stop the deployment by clicking **Abort**.

VxFlex OS components are installed on the nodes. Progress is displayed in the **Status** column.

6. Add SDS nodes to the system, and assign to Protection Domains (PD) and Fault Sets (FS):

By default, all nodes are added to the default PD.

a. In the **Add Nodes** screen, configure new PD and FS, according to the following table.

To do this	Perform this action
Add a PD	Click .
Add a FS to a PD	Select the PD, then click .
Assign a node to a PD or to a FS	Select a PD or FS, select the node, then click .

b. Click **Proceed**.

SDSs are added to the system, as configured. Progress is displayed in the **Status** column.

7. Configure caching for SSD and NVMe SSD devices:

a. In the **Configure Cache** screen, for each displayed SSD device, select whether to use it for storage or to provide caching for HDD devices on that node.

b. If using the device for caching, select the caching method (DAS Cache or RFcache). Available options differ according to the node configuration.

c. You can use these quick configuration options:

To do this	Perform this action
Apply the node's basic caching configuration (if the configuration exists)	Click **Basic caching configuration**.
Apply the current cache settings to all nodes of the same type in the cluster	Click **Apply to all nodes**.
Configure all SSD devices for storage use	Click **Set all as storage**.

d. Click **Apply**.

Node caching selections are configured. Progress is displayed in the **Status** column.

8. Add devices to SDS nodes.

   Every installed device must be added to a node.

   a. In the **Add Devices** screen, review or change the devices Storage Pool assignment.

   b. You can use these configuration options:

To do this	Perform this action
Create a new Storage Pool	Click and define the Storage Pool.
Assign a device to a Storage Pool	Select the device, select the Storage Pool, then click .
Move a device from one Storage Pool to another	Assign the device to the new Storage Pool.

   c. Click **Apply**.

   Storage Pools are configured. Progress is displayed in the **Status** column.

Results

The nodes have been added to the VxFlex Ready Node system.

After you finish

Note

A Red Hat license is required for storage-only configurations with Red Hat OS image downloaded from the product download 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.

The final screen of the deployment wizard enables you to perform the most common post-deployment tasks. Some tasks are mandatory, while others are recommended or optional. See the Post-deployment checklist on page 94.

Configuring the number of devices in a Storage Pool

Manually configure the number of devices in a Storage Pool by updating the `gatewayUser.properties` file.

By default, there can be a maximum of 300 devices in a Storage Pool. During deployment, the Storage Pools are populated with devices one by one such that in rare instances, the last Storage Pool can be left with devices from less than the minimum number of SDSs allowed. To proactively prevent this from occurring, you can change the number of devices allowed in a Storage Pool so that there will be more devices in the last Storage Pool. (Alternatively, post-deployment you can manually remove a device from a Storage Pool that has more than the minimum number of devices and add it to the Storage Pool with fewer devices than the minimum.)

Note

To ensure that this configuration change is applied during deployment, you must perform this procedure before the start of the deployment process.

Procedure

1. Calculate how many devices should be added to each Storage Pool so that the last Storage Pool will not have less than the minimum number of devices.
2. Using a text editor, open the `gatewayUser.properties` file, located in the following directory on the AMS server:
   - Windows 64-bit: `C:\Program Files\EMC\ScaleIO\AMS\webapps\ROOT\WEB-INF\classes\`
   - Linux: `/opt/emc/scaleio/ams/webapps/ROOT/WEB-INF/classes`
3. Edit the following parameter:

   `ams.extend_process_automation.max_devices_per_sp=`*x* where *x* is the maximum number of devices per Storage Pool, as calculated above.
4. Save and close the file.
5. Restart the ScaleIO AMS service:
   - Windows: Restart the EMC ScaleIO AMS service
   - Linux: Type the following command: `service scaleio-ams restart`

Extend a system with nodes of a different operating system

Add nodes to an existing VxFlex Ready Node system whose operating system is different than that of the existing system.

You can add VMware nodes to a Linux-based system, or Linux nodes to a VMware-based system.

Extend a VMware-based system with non-ESXi nodes

You can add non-ESXi nodes to a VMware-based VxFlex Ready Node system.

VMware-based VxFlex Ready Node systems are deployed (by default) as hyperconverged systems. You can add storage-only or VxFlex OS converged nodes of a different operating system.
If the current IP address range is sufficient, add the new nodes as normal, as described in Deploy or extend a VxFlex Ready Node system on page 85.

If the current IP address range needs to be modified, follow the steps in this procedure.

Procedure

1. From the **System Settings** > **Management and Data IPs** tab, in the **Management Network** panel:
   a. Select **Disable range validation (Mixed mode)**.
   b. Type the IP address ranges:
      - Host and iDRAC ranges are required only if the current range is not sufficient.

The IP address ranges do not need to match the range of the current system.

2. From the **Data Network** panel, type the IP address range for the new nodes:
   - Host IP address range is required only if the current range is not sufficient.

   The IP address ranges do not need to match the range of the current system.

3. Click **Save**, then close the window.

Results

The current system is configured to add nodes of the different operating system type.

After you finish

Extend the system as described in Deploy or extend a VxFlex Ready Node system on page 85.

Extend a Linux-based system with non-Linux nodes

You can add non-Linux nodes to a Linux-based VxFlex Ready Node system.

Linux-based VxFlex Ready Node systems are deployed (by default) as non-converged systems. You can add VxFlex OS converged nodes of a different operating system by following the procedure described in this topic, and then running the procedure in Deploy or extend a VxFlex Ready Node system on page 85.

Procedure

1. From the **System Settings** > **System** tab, select **Virtual host (ESX)**, then click **Save**.

   A warning is displayed notifying you of the upcoming necessary tasks.

2. From the **System Settings** > **Security** tab, type an SVM password, then click **Save**.
3. From the **System Settings** > **Management and Data IPs** tab, in the **Management Network** panel:

   a. Select **Disable range validation (Mixed mode)**.

   b. Type the IP address ranges:
      - SVM IP address range is mandatory.
      - Host and iDRAC ranges are required only if the current range is not sufficient.

      The IP address ranges do not need to match the range of the current system.

4. From the **Data Network** panel, type the IP address range for the new nodes:
   - SVM IP address range is mandatory.
   - Host IP address range is required only if the current range is not sufficient.

   The IP address ranges do not need to match the range of the current system.

5. Click **Save**, then close the window.

Results

The current system is configured to add nodes of the different operating system type.

After you finish

Extend the system as described in Deploy or extend a VxFlex Ready Node system on page 85.

CHAPTER 4

Post-Deployment Activities

This section describes activities that are performed after deployment of a VxFlex Ready Node system. Some activities are mandatory, while others are recommended or optional.

- Post-deployment checklist 94
- Create and map volumes 94
- Update the system component performance profiles 96
- Configuring ESRS 97
- Configure native users 98
- Configure LDAP users 99
- Configuring SNMP using the VxFlex OS GUI 100
- Setting credentials for existing Monitor users 100
- Back up SSH public keys using the VxFlex OS GUI 101
- Configure the timing of the alert on certificate expiration 102
- Sign and update AMS security certificates 102
- Activating the ESX license on a node 103
- Adding an external SDC to an existing VxFlex OS system 103
- Uninstall VxFlex Ready Node components 105
- Prevent certificate errors when ESX hosts are added using hostnames 107

Post-deployment checklist

Use the following checklist to verify that you complete the required and optional post-deployment procedures.

Post-deployment procedures

Procedure	Notes	Mandatory?	Done
Add and map volumes	Adding and mapping volumes can be performed using several management tools. For more information, see Create and map volumes on page 94	Mandatory	
Update the system component performance profile	For more information, see Update the system component performance profiles on page 96.	Mandatory	
Install VxFlex OS license	A valid VxFlex OS license must be installed. For information how to obtain and install the license, see Licensing on page 109.	Mandatory	
Review VxFlex OS performance best practices		Recommended	
Review VxFlex OS security best practice configurations		Recommended	
Configure EMC Secure Remote Support (ESRS)	For more information, see Configuring ESRS on page 97.	Optional	
Define native users in the system and change the default admin password.	For instructions, see Configure native users on page 98.	Optional	
Define LDAP users (if using VxFlex OS with LDAP).	For instructions, see Configure LDAP users on page 99. For detailed information, see *User Roles and LDAP Usage Technical Notes*.	Optional	
Enable SNMP and configure the SNMP trap receiver	For more information, see Configuring SNMP using the VxFlex OS GUI on page 100.		
Change monitor user password		Optional	
Back up all SSH keys from nodes and Storage VMs		Optional	

Create and map volumes

Create volumes from devices added to SDS nodes, and then map the volumes to SDC nodes.

You can create and map volumes using various management tools, including the CLI, VxFlex OS GUI, and REST API.

Creating and mapping volumes is necessary 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. For additional details on creating and mapping volumes, see the user documentation.

Creating and mapping volumes using the VxFlex OS GUI

Use the VxFlex OS GUI to create a volume and map it to an SDC.

Before you begin

Ensure that you can authenticate to the VxFlex OS GUI with an admin user.

For more information on creating and mapping volumes using the VxFlex OS GUI, see the user documentation.

Procedure

1. Open the VxFlex OS GUI and log in with credentials:
   - Linux: Run the script `/opt/emc/scaleio/gui/run.sh`.
   - Windows: Click **Start** > **All Programs** > **VxFlex OS GUI**

   The VxFlex OS GUI is displayed.
2. Create volumes:
   a. In the **Frontend** > **Volumes** view, select a Storage Pool to which to add the volume.
   b. From the **Command** menu or context-sensitive menu, select **Add Volume**.

      The **Add Volume** window is displayed.
   c. To create more than one volume, select **Create multiple volumes** and type the number of volumes in the **Copies** box.
   d. Type a name in the **Name** box.
   e. 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.
   f. Type a number in the **Size** box, representing the volume size in GB (allocation granularity is 8 GB).
   g. Select **Thick** (default) or **Thin** provisioning options.
   h. Enable or disable the RAM Read Cache feature by selecting or clearing **Use RAM Read Cache**.
   i. Click **OK**.
3. Map volumes to an SDC:
   a. In the **Frontend** > **Volumes** view, select the volumes.
   b. 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.
   c. In the **Select Nodes** panel, select one or more SDCs to which you want to map the volumes.
   d. 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.

Creating and mapping volumes using the CLI

Create a volume and map a volume to an SDC using the CLI.

Use the following SCLI commands to create and map volumes. For more details on SCLI command usage, see the *VxFlex Ready Node CLI Reference Guide*.

Procedure

1. Create a volume:

```
scli --add_volume --size_gb <SIZE> --volume_name <VOL_NAME> --
protection_domain_name <PD_NAME> --storage_pool_name
<SP_NAME> --mdm_port 6611
```

2. Map a volume to an SDC:

```
scli --map_volume_to_sdc --volume_name <VOL_NAME> --sdc_ip
<SDC_IP> --mdm_port 6611
```

Update the system component performance profiles

Set the performance profile of all MDM, SDC, and SDS components to High.

Before you begin

Ensure that you can authenticate to the VxFlex OS GUI or the CLI with an admin user.

You can use the VxFlex OS GUI to change each performance profile separately, as described. Alternatively, you can log in to the CLI and set all parameters at once by running this command: `scli --set_performance_parameters --all_sds --all_sdc --apply_to_mdm --profile high_performance`

Procedure

1. Log in to the VxFlex OS GUI.
2. Set performance profile for all MDMs:
   a. From the **Backend > Storage** view, select the System object.
   b. From the **Command** menu or context-sensitive menu, select **Set Performance Profile for MDMs**.
   c. In the **Set Performance Profile for** dialog, select **High**, and click **OK**.
3. Set performance profile for all SDSs:
   a. From the **Backend > Storage** view, select the System object.
   b. From the **Command** menu or context-sensitive menu, select **Set Performance Profile for all SDSs**.
   c. In the **Set Performance Profile for** dialog, select **High**, and click **OK**.
4. Set performance profile for all SDCs:

a. From the **Frontend > SDCs** view, select the System object.

b. From the **Command** menu or context-sensitive menu, select **Set Performance Profile for all SDCs**.

c. In the **Set Performance Profile for** dialog, select **High**, and click **OK**.

Results

The performance profiles for all components are set to High.

Configuring ESRS

Configure EMC Secure Remote Support (ESRS) for remote support.

Before you begin

Ensure you have:

- Authentication credentials for an admin user for the VxFlex OS GUI.
- One or more IP addresses of the ESRS servers. Note that ESRS does not currently support IPv6.
- ESRS username and password.
- 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).
- 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.

Procedure

1. Log in to the VxFlex OS GUI.
2. From the **Preferences** menu at the top right side of the window, select **System Settings** > **ESRS**.

3. In the **Set ESRS Server** window, add the first ESRS server by typing its IP address in the **ESRS IP addresses** box.

   To add another ESRS server, click the **Add** button .
4. In the **User Name** box, type the ESRS username.
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 not want to use the ESRS feature, and you want to disable the alert for unregistered ESRS gateway, select **Allow configuration with no ESRS registered**.
8. Click **Save**.

9. Enable ESRS registration to Gateway, using one of the following methods:

Option	Description
Use the FSOGWTool to add the ESRS-GW details to a DNS server known by the Gateway node.	Run this command on the Gateway: `/opt/emc/scaleio/gateway/bin/FSOGWTool.sh --register_esrs_gateway --scaleio_gateway_ip 10.234.221.46 --scaleio_gateway_user admin --scaleio_gateway_password 123 --esrs_gateway_ip ESRS-VE-2 --esrs_gateway_user Scott.jones@nordstrom.com --esrs_gateway_password Password1 --connect_in_ip 10.234.221.46` Where --esrs_gateway_ip is the hostname of the ESRS Gateway
Modify the /etc/hosts file on the Gateway	Add a line to /etc/hosts that identifies the IP address with the hostname of the ESRS Gateway, for example: 10.234.223.220 ESRS-VE-2

Configure native users

Configure native users using the VxFlex OS CLI.

Use the following SCLI commands to create and modify native VxFlex OS users. 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*. 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
```

Procedure

1. Log in as an admin user to the CLI:

   ```
   scli --login --username <NAME> --password <PASSWORD>
   ```

   where *<PASSWORD>* is the password used for the first login.

2. (Optional) Change the default admin user password:

   ```
   scli --set_password --old_password <PASSWORD> --new_password
   <NEW_PASSWORD>
   ```

3. Add users:

```
scli --add_user --username <NAME> --user_role <ROLE --
mdm_port 6611
```

A randomly generated password for the created user is returned.

4. (Optional) Log in with the new user and then change its password:

```
scli --login --username <NAME> --password <PASSWORD>
```

```
scli --set_password --old_password <OLD_PASSWORD> --
new_password <NEW_PASSWORD> --mdm_port 6611
```

This is optional because the VxFlex OS GUI, CLI, or REST will enforce a password change if user logs in with the original password.

Configure LDAP users

Configure LDAP users if you are using LDAP with VxFlex OS.

For more information on using LDAP with VxFlex OS, see the *VxFlex OS User Roles and LDAP Usage Technical Notes*. To run SCLI commands, you must be logged in. Actual command syntax is operating-system dependent. For more information, see the *VxFlex OS CLI Reference Guide*.

Procedure

1. Create Active Directory (AD) groups that correspond to the user roles offered by VxFlex OS.
2. Use the CLI to add LDAP service to the MDM:

```
scli --add_ldap_service --ldap_service_uri <URI> --
ldap_base_dn <LDAP_DN> --mdm_port 6611
```

Note

VxFlex OS systems support authentication by up to eight LDAP servers. When multiple LDAP servers are used, add each one separately using this command.

---

3. Use the CLI to assign an LDAP group to the user role:

```
scli --assign_ldap_groups_to_roles (--ldap_service_id
<LDAP_SERVICE_ID> | --ldap_service_name <LDAP_SERVICE_NAME>)
[--administrator_role_dn <LDAP_GROUP_DN>] [--security_role_dn
<LDAP_GROUP_DN>] [--backend_config_role_dn <LDAP_GROUP_DN>]
[--frontend_config_role_dn <LDAP_GROUP_DN>] [--
monitor_role_dn <LDAP_GROUP_DN>] --mdm_port 6611
```

4. Use the CLI to set the system to mixed authentication method, LDAP and native:

```
scli --set_user_authentication_method --
native_and_ldap_authentication --mdm_port 6611
```

5. Log in again to apply these changes.

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.

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 13 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.

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.

Configure the timing of the alert on certificate expiration

Configure the number of days before certificate expiration when the alert is displayed.

The system's default alert that its certificate will expire is displayed 30 days before the expiration date. You can edit the properties file to configure this property.

Procedure

1. Using a text editor, open the `gatewayUser.properties` file, located in the following directory on the AMS server:
   - Windows 64-bit: `C:\Program Files\EMC\ScaleIO\AMS\webapps\ROOT\WEB-INF\classes\`
   - Linux: `/opt/emc/scaleio/ams/webapps/ROOT/WEB-INF/classes`
2. Locate the line in the file that contains ams.certificateAboutToExpireAlertTimeInDays.
3. Edit the value of the ams.certificateAboutToExpireAlertTimeInDays.
4. Save and close the file.
5. Restart the ScaleIO AMS service:
   - Windows: Restart the EMC ScaleIO AMS service
   - Linux: Type the following command: `service scaleio-ams restart`

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.

Activating the ESX license on a node

Activate the ESX license on a node, after its deployment.

The ESX image does not include an ESX license. Therefore, once the node deployment is complete, it is recommended that you install the ESX license on the node, or else, it will expire after 60 days.

For information on how to activate the ESX license, click https://kb.vmware.com/kb/2014295.

Adding an external SDC to an existing VxFlex OS system

In a 2-layer architecture, after deploying the system, install the SDC according to the operating system-specific instructions in the following section. The SDCs will be connected to the existing system.

Install the SDC on an ESX server and connect it to VxFlex OS

Install the SDC with the appropriate parameters to connect it to an existing VxFlex OS system.

Before you begin

Ensure that you have:

- The virtual IP address or MDM IP address of the existing system
- Login credentials for the SDC
- The appropriate installation packages for the SDC
- Access to the drv_cfg tool. Contact Customer Support for access to this tool on ESX.

The following procedure describes installing an external SDC on an ESX server using the esxcli. Alternatively, you can install the external SDC using the vSphere plug-in. For more information, see "Installing the SDC on ESX hosts" in the *VxFlex OS Deployment Guide*.

Procedure

1. On the ESX on which you are installing the SDC, set the acceptance level:

```
esxcli --server=<SERVER_NAME> software acceptance set --
level=PartnerSupported
```

   where *<SERVER_NAME>* is the ESX on which you are installing the SDC.

2. Install the SDC:

```
esxcli software vib update -d “Full Path”
```

3. Set the IP address of the MDM:

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

   where

   - *<LIST_VIP_MDM_IPS>* is a comma-separated list of the MDM IP addresses or the virtual IP address of the MDM
   - *<XXXXXX>* is the version

Results

The SDC is installed on the ESX server and is connected to VxFlex OS.

Installing SDC on a Linux server and connecting it to VxFlex OS

Install the SDC with the appropriate parameters to connect it to an existing VxFlex OS system.

Before you begin

Ensure that you have:

- The virtual IP address or MDM IP address of the existing system
- Login credentials for the SDC
- The appropriate installation packages for the SDC

The following procedure describes manually installing an external SDC on a Linux server. On most servers (with the exception of hLinux), you can install the external SDC using the VxFlex OS Installation Manager. For more information, see "Extending an existing VxFlex OS system" in the *VxFlex OS Deployment Guide*.

Note

External SDC on RHEL 7.4 is supported on bare-metal servers only, not as guests on a hypervisor.

Procedure

1. Install the SDC:
   - RHEL/CentOS /Oracle Linux

```
MDM_IP=<LIST_VIP_MDM_IPS> rpm -i <SDC_PATH>.rpm
```

- Ubuntu. These files must be extracted before use, as described in the "Manual deployment on physical servers" section of the *VxFlex OS Deployment Guide*.

```
MDM_IP=<LIST_VIP_MDM_IPS> dpkg -i <SDC_PATH>.deb
```

- CoreOS

```
MDM_IP=<LIST_VIP_MDM_IPS> ./<LIST_VIP_MDM_IPS>.bsx
```

where

- *<LIST_VIP_MDM_IPS>* is a comma-separated list of the MDM IP addresses or the virtual IP address of the MDM
- *<SDC_PATH>* is the path where the SDC installation package is located

Results

The SDC is installed on the Linux server and is connected to VxFlex OS.

Installing SDC on a Windows server and connecting it to VxFlex OS

Install the SDC with the appropriate parameters to connect it to an existing VxFlex OS system.

Before you begin

Ensure that you have:

- The virtual IP address or MDM IP address of the existing system
- Login credentials for the SDC
- The appropriate installation packages for the SDC

The following procedure describes manually installing an external SDC on a Windows server. Alternatively, you can install the external SDC using the VxFlex OS Installation Manager. For more information, see "Extending an existing VxFlex OS system" in the *VxFlex OS Deployment Guide*.

Procedure

1. On the Windows server on which you are installing the SDC, run:

```
msiexec /i <SDC_PATH>.msi MDM_IP=<LIST_VIP_MDM_IPS>
```

where

- *<SDC_PATH>* is the path where the SDC installation package is located
- *<LIST_VIP_MDM_IPS>* is a comma-separated list of the MDM IP addresses or the virtual IP address of the MDM

Results

The SDC is installed on the Windows server and is connected to VxFlex OS.

Uninstall VxFlex Ready Node components

Uninstall the VxFlex Ready Node AMS or GUI.

Uninstalling the VxFlex Ready Node AMS from a Windows server

Uninstall the AMS from a Windows server.

Procedure

1. Uninstall the AMS using one of the following methods:
   - AMS installer:
     a. Run the AMS install wizard.
     b. Select the option to remove the program.
     c. Proceed through the rest of the wizard.
   - Windows Control Panel:
     a. From Windows Start, select **Control Panel** > **Programs and Features**.
     b. Select **VxFlex OS AMS** from the list of programs and click **Uninstall**.
2. Reboot the server.

Uninstalling the VxFlex Ready Node AMS from a Linux server

Uninstall the AMS when using Linux.

Procedure

1. Log in to the AMS server.
2. Run the following command:

```
rpm -e EMC-ScaleIO-ams-2.6-<build>.x86_64
```

Note

During the uninstallation process, the following warnings may be displayed. They are not an indication of a problem and can be ignored:

```
warning: file /opt/emc/scaleio/ams/scli/scli-sles12-1: remove failed:
No such file or directory
warning: file /opt/emc/scaleio/ams/scli/scli-sles12: remove failed:
No such file or directory
warning: file /opt/emc/scaleio/ams/scli/scli-rhel7: remove failed: No
such file or directory
warning: file /opt/emc/scaleio/ams/scli/scli-rhel6: remove failed: No
such file or directory
```

---

3. Reboot the server.

Uninstall the VxFlex OS GUI

Uninstall the VxFlex OS GUI using the installer or Windows Control Panel.

Procedure

1. Uninstall the VxFlex OS GUI using one of the following methods:
   - Using the installer:
     a. Run the VxFlex OS GUI install wizard.

b. Select the option to remove the program.

c. Proceed through the rest of the wizard.

- Using the Windows Control Panel:

  a. From Windows Start, select **Control Panel** > **Programs and Features**.

  b. Select **VxFlex OS GUI** from the list of programs and click **Uninstall**.

Prevent certificate errors when ESX hosts are added using hostnames

When hosts are added to the vCenter using hostnames, a certificate alert raised by the AMS, and certificate renewal fails. Edit the `ams.properties` file to prevent occurrence of this issue.

Note

The change made to the file that is edited in this workaround is not saved during upgrades, and may need to be performed again.

Procedure

1. On the host where AMS is installed, open the following file in a text editor:

Operating system	File path
Linux	`/opt/emc/scaleio/ams/webapps/ROOT/WEB-INF/classes/ams.properties`
Windows	`C:\Program Files\EMC\scaleio\AMS\webapps\ROOT\WEB-INF\classes\ams.properties`

2. Change the property `verifyEsxCertificates=True` to `verifyEsxCertificates=False`.
3. Save the `ams.properties` file.
4. Restart the VxFlex Ready Node AMS service:

Operating system	Action
Linux	Type the following command: `service scaleio-ams restart`
Windows	Restart the EMC ScaleIO AMS service.

CHAPTER 5

Licensing

This section contains procedures on how to obtain and activate the electronic license for your VxFlex Ready Node software.

- Licensing overview 110
- Activating entitlements and installing a license file 111
- License file example 116
- Error messages 116

Licensing overview

VxFlex OS licenses are purchased by physical device capacity (in TB). You can activate your licensed capacity over multiple VxFlex OS systems—each system with its unique installation ID.

You download VxFlex OS licenses from the EMC Software Licensing Central website, using the procedures described in "Activating entitlements and installing a license file". Then, you install the licenses on VxFlex OS, as described in Install the license - using the VxFlex OS GUI on page 115.

You can view current license information using the CLI or the VxFlex OS GUI.

The following steps summarize the licensing process:

1. Purchase VxFlex OS, and receive a License Authorization Code (LAC) email with a link to the licensing site.

   If you do not have the LAC email, you can search for the LAC number from the EMC Software Licensing Central website, by entering the Sales Order number and using the **Search Entitlements** option.

2. Retrieve the installation ID from VxFlex OS.
3. Click the link in the LAC email, and use the online wizard to complete the entitlement activation process.

   a. Save the license file, and install it using the CLI.

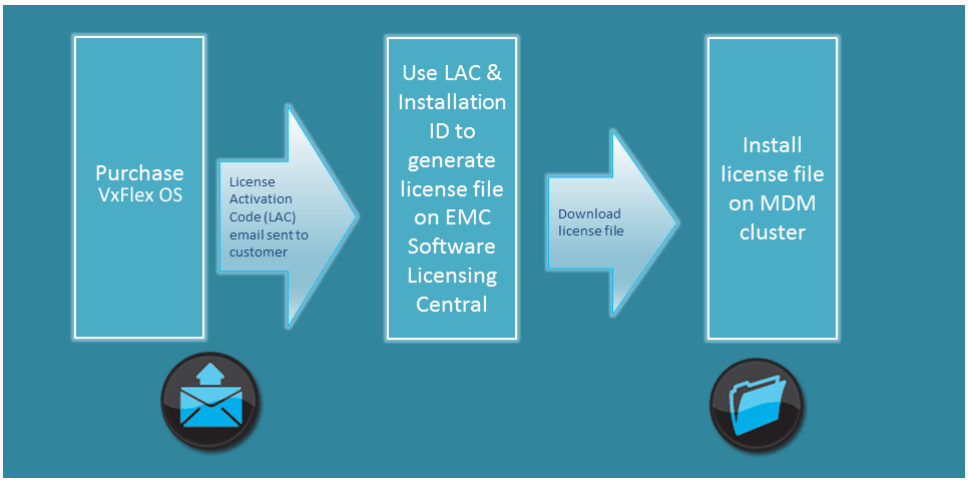

The following table describes VxFlex OS eLicensing terminology.

Table 25 eLicensing terminology(continued)

Term	Description
Dell EMC Online Support	The online support portal, http://support.emc.com, contains product support information and links to the Software Licensing Central web site.
Entitlements	The Dell EMC Software Licensing Central web site lists the entitlements (usage rights) that you have

Table 25 eLicensing terminology(continued) (continued)

Term	Description
	purchased, that you can activate for a specific host machine.
LAC email	Email sent to a customer who has purchased an Dell EMC product, containing a License Authorization Code (LAC), which is needed to complete the entitlement activation process on the EMC Software Licensing Central web site.

When you purchase a license entitlement for VxFlex OS, a License Authorization Code (LAC) email is sent to you, or to your purchasing department. If you cannot find the LAC email, you can use the Software Licensing Central website to find your license entitlements.

The following figure shows a sample LAC email:

Figure 14 Licensing LAC email

Dear EMC Software User,
Thank you for choosing EMC software. Your EMC Software License Authorization Code (LAC) is [redacted].
You must redeem this LAC for license keys to activate your software. Please protect your LAC like you would any other license key to prevent anyone from improperly activating your software.

Activating Your Software
1. Click here or copy and paste the following URL (https:[redacted]
LAC=[redacted]) into a web browser to activate your entitlements.
2. You will be prompted to log in. (New users should follow the new member registration steps).
3. Follow the on-screen instructions.

Downloading Your Software
1. Click here or copy and paste the following URL (https://ngtest-ci.emc.com/downloads/) into a web browser to download your software.
2. You will be prompted to log into EMC's Online Download Service Center (New users should follow the new member registration steps).
3. Enter the product name in the search field to find the software you wish to download.

License Authorization Code: RR91DYQLF9RTXZT9RMQH

Product #	Title	Quantity
456-106-154	EMC SCALEIO SOFTWARE CAPACITY=CB	250
456-106-155	EMC SCALEIO ENTERPRISE FEATURES=CB	250

If you have any questions about your sales order please contact your EMC Account Representative or your Authorized Reseller.

Activating entitlements and installing a license file

VxFlex OS licenses are assigned to VxFlex OS systems, each of which is identified by a unique installation ID.

You use the VxFlex OS installation ID, together with your LAC, to activate the entitlement and then download the license file. Then, you install this file in your MDM cluster.

Activate an entitlement and download the license file

This section describes how to activate the entitlement that was purchased.

VxFlex OS is procured by total capacity, but you can activate portions of this total capacity over multiple VxFlex OS installations. For example, your purchase order may have been for 1000 TB. Your LAC will entitle you to activate all, or part of that. You can activate 500 TB for one VxFlex OS installation, and leave the rest for another activation, for the same, or a different system.

To activate the entitlement, perform the following:

Procedure

1. Identify the VxFlex OS installation ID:
   - Using the CLI:
     Run the following command:

     ```
     scli --query_license
     ```

     The installation ID is displayed:

     ```
     Installation ID: 0123456789abcdef
     ```

     Note

     To run CLI commands, you first need to log in. For more information see the "Logging In" section in the user documentation.

     Actual command syntax is operating-system dependent. For more information, see the *VxFlex Ready NodeCLI Reference Guide*.

   - Using the VxFlex OS GUI:
     From the top right of the main window, open the drop-down menu that appears next to the user name, and select **About**.

     The installation ID is displayed in the **About** window.
2. If you have the LAC email, skip to step 4 .
3. If you do not have your LAC email, perform the following:
   a. From the Dell EMC support website, browse to the Software Licensing Central system:
      a. Open the support website: http://support.emc.com.

         If you are a new user, create a new user account.
      b. From the **Support Tasks** list, click **Manage Licenses and Usage Intelligence**.
      c. From the software list, click **VxFlex OS** . The **Powerlink Licensing** website is displayed.
      d. Click **View Entitlements**. The **Search Entitlements** screen appears.

   b. Type the Sales Order number, then click **Search Entitlements**.

      A list of entitlements is displayed.

   c. Locate the entitlement to activate, and choose **Options** > **Activate**.

   The **Powerlink Licensing—Search Entitlements to Activate** screen appears. Skip to step 6.
4. If you have your LAC email, perform the following:
   a. Click the link in the LAC email, and log in.

      The **Activate—Search for Products** screen appears:

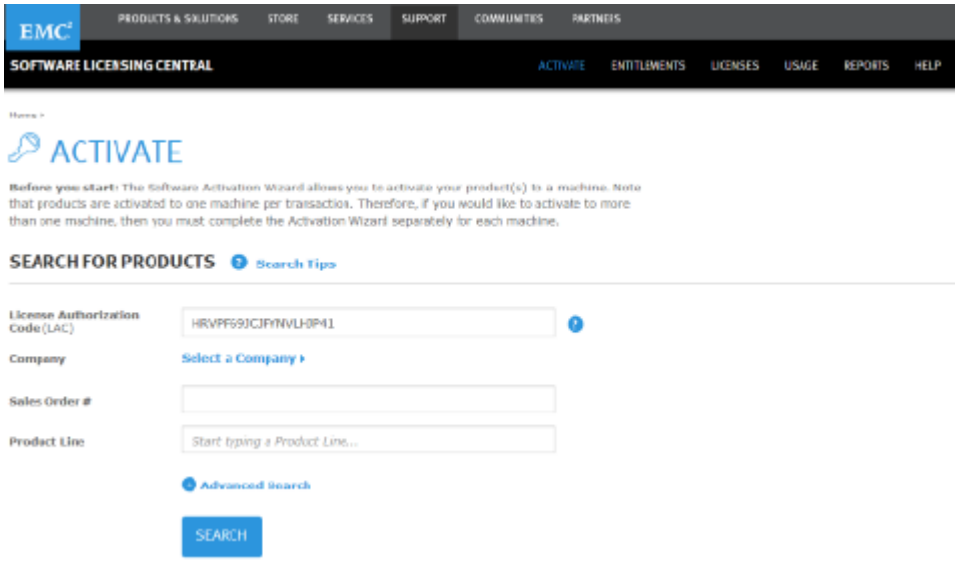

b. Enter your LAC code, or search by Sales Order number, then click **Search**.

The **Select Products** screen appears:

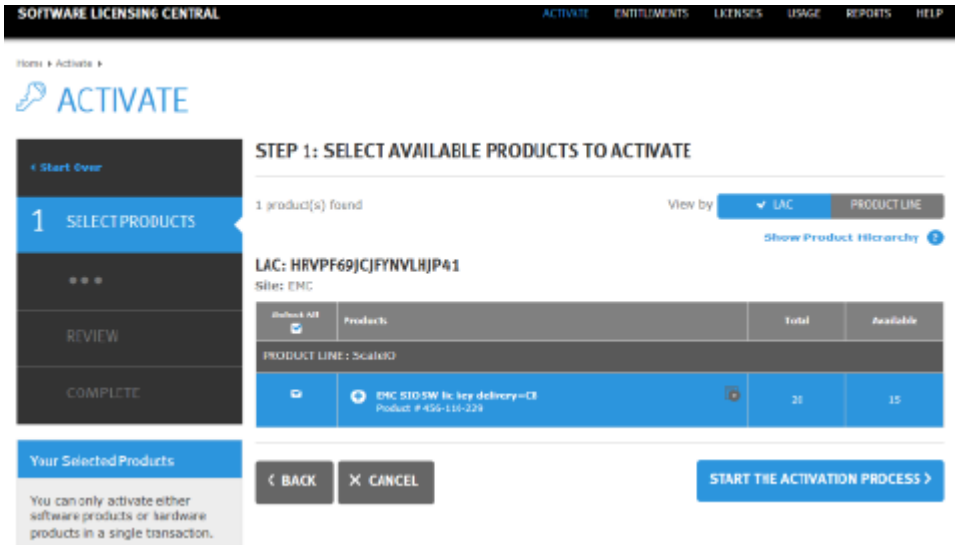

5. In the **Select Products** screen, select the product to activate, and click **Start the Activation Process**.
6. In the **Company Details** screen, confirm (or update) company information, and click **Select a Machine**.

   The **Select a Machine** screen appears:

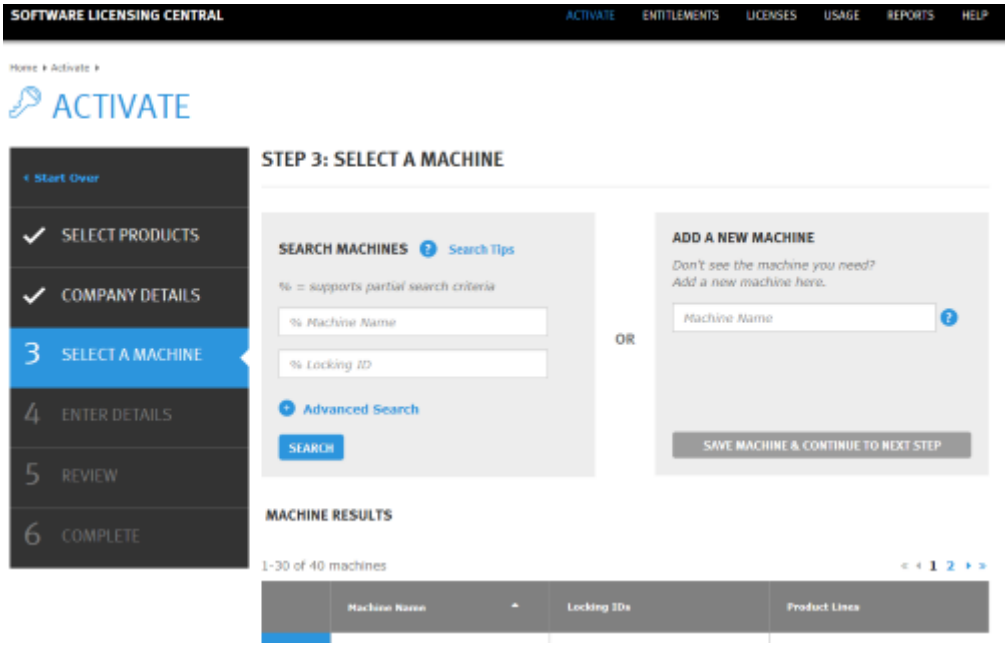

7. In the **Select a Machine** screen, select a machine on which to activate the product in one of these ways:
   - Click **Search** to locate an existing machine (one on which EMC product was previously activated).
   - Add a new machine name, then click **Save Machine & Continue**.

   In the context of the activating process, a machine is a VxFlex OS system, which could comprise multiple servers.

   The **Enter Details** screen appears:

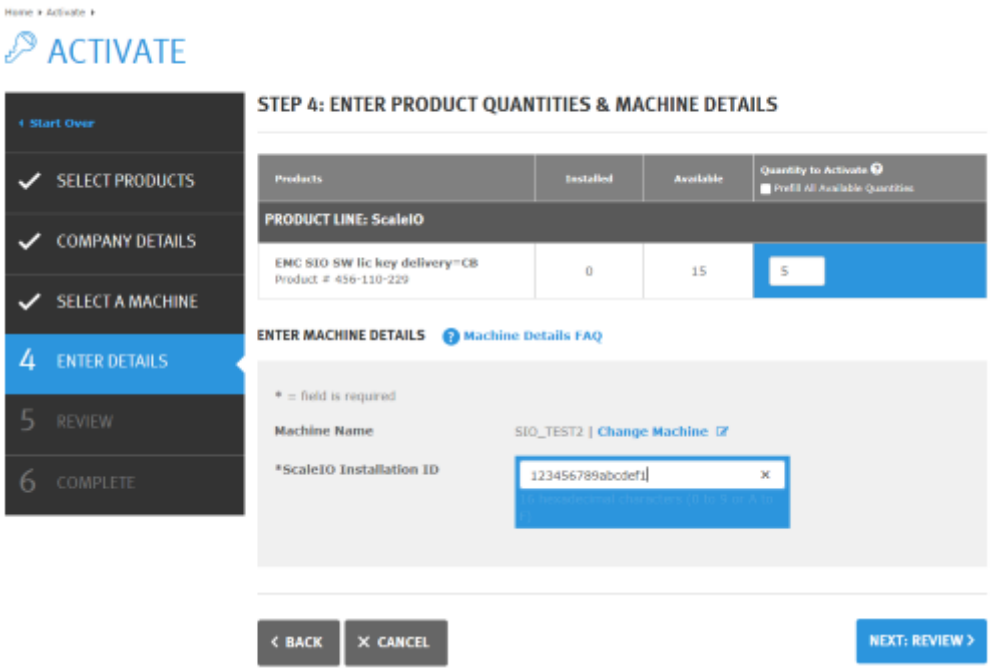

8. In the **Enter Details** screen, enter the following:
   - **Quantity** (in TB) to activate on this machine.

To allocate the available capacity over multiple machines, select less than the full amount available, and repeat the activation process on the other machine.

- VxFlex OS Installation ID, from the beginning of this procedure.

9. Click **Next**.
10. In the **Review** screen, you can review your selections.

    The license key will be emailed to the username that is logged in to the licensing system. To send it to more recipients, click **Email to more people** and enter their email addresses.

11. Click **Activate**.

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.

Install the license - SCLI

To install the license, run the following command:

```
scli --set_license --license_file <license_file>
```

where *<license_file>* is the full path to the license file

Example: `scli --set_license --license_file /tmp/0239SH4SS89023T6.lic`

The VxFlex OS license is now installed on the MDM cluster.

You can view license information using the `query_license` command and from the **About** menu in the VxFlex OS GUI.

Actual command syntax is operating-system dependent. For more information, see the *VxFlex Ready Node CLI Reference Guide*.

License file example

The following figure illustrates a license file with a license for 200TB of capacity:

Figure 15 License file example

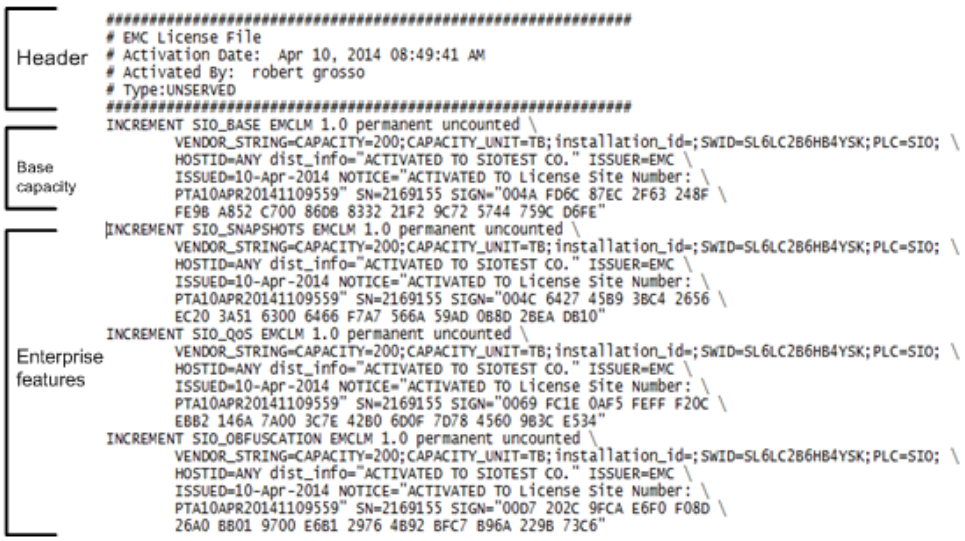


```
#############################################################
# EMC License File
# Activation Date:  Apr 10, 2014 08:49:41 AM
# Activated By:  robert grosso
# Type:UNSERVED
#############################################################
INCREMENT SIO_BASE EMCLM 1.0 permanent uncounted \
        VENDOR_STRING=CAPACITY=200;CAPACITY_UNIT=TB;installation_id=;SWID=SL6LC2B6HB4YSK;PLC=SIO; \
        HOSTID=ANY dist_info="ACTIVATED TO SIOTEST CO." ISSUER=EMC \
        ISSUED=10-Apr-2014 NOTICE="ACTIVATED TO License Site Number: \
        PTA10APR20141109559" SN=2169155 SIGN="004A FD6C 87EC 2F63 248F \
        FE9B A852 C700 86DB 8332 21F2 9C72 5744 759C D6FE"
INCREMENT SIO_SNAPSHOTS EMCLM 1.0 permanent uncounted \
        VENDOR_STRING=CAPACITY=200;CAPACITY_UNIT=TB;installation_id=;SWID=SL6LC2B6HB4YSK;PLC=SIO; \
        HOSTID=ANY dist_info="ACTIVATED TO SIOTEST CO." ISSUER=EMC \
        ISSUED=10-Apr-2014 NOTICE="ACTIVATED TO License Site Number: \
        PTA10APR20141109559" SN=2169155 SIGN="004C 6427 45B9 3BC4 2656 \
        EC20 3A51 6300 6466 F7A7 566A 59AD 0B8D 2BEA DB10"
INCREMENT SIO_QoS EMCLM 1.0 permanent uncounted \
        VENDOR_STRING=CAPACITY=200;CAPACITY_UNIT=TB;installation_id=;SWID=SL6LC2B6HB4YSK;PLC=SIO; \
        HOSTID=ANY dist_info="ACTIVATED TO SIOTEST CO." ISSUER=EMC \
        ISSUED=10-Apr-2014 NOTICE="ACTIVATED TO License Site Number: \
        PTA10APR20141109559" SN=2169155 SIGN="0069 FC1E 0AF5 FEFF F20C \
        EBB2 146A 7A00 3C7E 42B0 6D0F 7D78 4560 9B3C E534"
INCREMENT SIO_OBFUSCATION EMCLM 1.0 permanent uncounted \
        VENDOR_STRING=CAPACITY=200;CAPACITY_UNIT=TB;installation_id=;SWID=SL6LC2B6HB4YSK;PLC=SIO; \
        HOSTID=ANY dist_info="ACTIVATED TO SIOTEST CO." ISSUER=EMC \
        ISSUED=10-Apr-2014 NOTICE="ACTIVATED TO License Site Number: \
        PTA10APR20141109559" SN=2169155 SIGN="00D7 202C 9FCA E6F0 F08D \
        26A0 BB01 9700 E6B1 2976 4B92 BFC7 B96A 229B 73C6"
```

The license file includes the following sections:

- Header: General information
- General license: Shows the capacity licensed for the system, in this case 200 (TB).

Error messages

The following table lists error messages that may be generated by the system and their troubleshooting solutions.

Table 26 Licensing error messages

Error Message	Description	Solution
The license key is invalid or does not match this version. Contact Support.	The license key is invalid.	Contact support.
The current system configuration exceeds the license entitlements.	More capacity has been installed than the license allows.	Reduce capacity, or extend the license capacity.

Table 26 Licensing error messages (continued)

Error Message	Description	Solution
Operation could not be completed. The license capacity has been exceeded.	When you try to add an SDS or device, it will cause the licensed capacity to be exceeded.	Do not add the SDS or device, or extend the license capacity.
The license key is too long	The license file is larger than expected.	Check the accuracy of the license key.
The license has expired	The duration of the license has ended.	Extend the duration of the license.
The license installation ID does not match the ID of this system	When the Installation ID was entered in the ELM, it may have been incorrect.	Contact support.
The license contains a mismatch of the SWID. Contact Support.	The license key is invalid.	Contact support.
The issuer of the license you are attempting to add does not match that of the product	The license key is invalid.	Contact support.
The license contains a mismatch of the capacity values for basic and advanced features. Contact Support.	The capacity licensed for basic features is not equal to the capacity licensed for advanced feature.	Contact support.

PART 3

Reference

This part provides additional information on advanced reference topics.

Chapter includes:

Chapter 6, "Common Tasks"

Chapter 7, "Server factory reset and power operations"

CHAPTER 6

Common Tasks

This section contains common task that are performed when working with VxFlex Ready Node

- Associating VxFlex OS volumes with physical disks 122
- Volume information - Linux 122
- Port usage and changing default ports 123
- Upgrading Java 124
- Move ESXi hosts to a new vCenter and reconfigure AMS 124
- Best practices for using VMware vMotion 125
- Update the hardware using the iDRAC virtual console 127
- AMS management takeover of non-managed nodes 129

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 27 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 27 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 16 Network ports

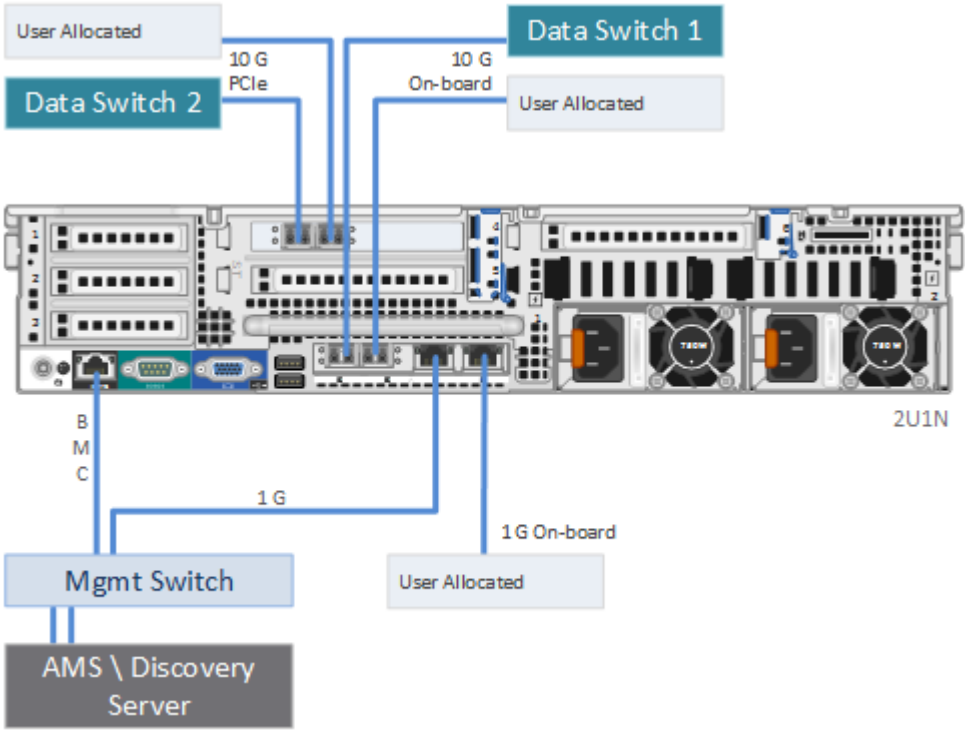

Disable VMware Admission Control

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 128.

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 127 to 6 on page 127 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 .

    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 142.

2. Boot the VxFlex Ready Node server to the Hardware ISO.

   For instructions, see Update the hardware using the iDRAC virtual console on page 127.

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 136.

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 129.
- 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 129.

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 7

Server factory reset and power operations

This section contains information on how to open a Dell iDRAC console session, perform power operations, and perform a factory reset.

- Perform a factory reset 138
- KVM activities 146

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 142.

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**.

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