

Dell EMC VxBlock™ System 1000

Architecture Overview

Document revision 1.14

December 2019

Copyright © 2018-2019 Dell Inc. All rights reserved.

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 Technologies, Dell, EMC, 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

Revision history	7
Introduction	9
System overview	10
Minimum and maximum configurations.....	11
Scale up compute resources.....	15
Network topology.....	15
Host boot configuration.....	18
Compute layer	19
Cisco UCS fabric interconnects.....	20
Cisco UCS servers.....	20
Compute server connectivity.....	23
Cisco UCS B-Series server connectivity.....	23
Cisco UCS C-Series Server direct connect.....	27
Cisco UCS C-Series Gen 3 Fabric Extender connection.....	29
Cisco UCS C-Series Gen 4 Fabric Extender connection.....	30
Mixed Cisco UCS B-Series and C-Series servers.....	31
Cisco Trusted Platform Module.....	33
Disjoint Layer 2 configuration.....	33
Bare metal support policy.....	35
Network layer	36
Data network routing topology.....	36
Management network layer.....	36
Cisco Nexus 31108TC-V Switch.....	37
Cisco Nexus 3232C Switch.....	37
Cisco Nexus 9336C-FX2 Switch.....	38
VxBlock System 1000 with no management aggregation layer.....	38
Single VxBlock System 1000 with a management aggregation layer.....	39
Multiple VxBlock Systems 1000 with a management aggregation layer.....	40
Software licenses.....	43
Cisco Nexus 93180YC-FX Switch.....	43
Cisco Nexus 93180LC-EX Switch.....	44
Cisco Nexus 9336C-FX2 Data Plane Switch (Cisco UCS Gen 3).....	46
Cisco Nexus 9336C-FX2 Switch (Cisco UCS Gen 4).....	47
Cisco Data Center Network Manager	49
Vscale Architecture and ACI readiness.....	49
VMware NSX factory integration.....	50
Storage network layer.....	50
Cisco MDS 9148S Multilayer Fabric Switch.....	52
Cisco MDS 9148T Multilayer Fabric Switch.....	52
Cisco MDS 9396S 16G Multilayer Fabric Switch	52
Cisco MDS 9396T 32G Multilayer Fabric Switch	53
Cisco MDS 9706 and 9710 Multilayer Directors.....	53
Storage layer	54
VMAX All-Flash storage overview.....	54
Array components and specifications.....	54
Storage features.....	55

Software licensing.....	56
VMAX storage scale.....	57
PowerMax storage overview.....	58
Array components and specifications.....	58
Storage features.....	60
Software licensing.....	61
PowerMax storage scale.....	62
Dell EMC Unity storage overview.....	62
Array components and specifications.....	63
Dell EMC Unity storage features.....	65
FAST Cache and FAST VP for Dell EMC Unity Hybrid arrays.....	66
Software licensing.....	67
Storage scale.....	68
Dell EMC Unity XT storage overview.....	68
Array components and specifications.....	69
Dell EMC Unity XT storage features.....	71
FAST Cache and FAST VP for Dell EMC Unity XT Hybrid arrays.....	72
Software licensing.....	72
Storage scale.....	73
Isilon storage overview.....	74
Array components and specifications.....	75
Storage features.....	76
Software licensing.....	76
Isilon storage scale.....	77
XtremIO storage overview.....	78
Array components and specifications.....	78
Storage features.....	79
XtremIO X2 storage scale.....	80
Supported Integrated Data Protection options.....	80
Multiple storage array overview.....	81
Block storage.....	81
File storage.....	81
Manage and monitor multiple storage arrays.....	81
Management.....	83
VxBlock Central options.....	83
VxBlock Central Base option.....	83
VxBlock Central Workflow Automation option.....	84
VxBlock Central Advanced Analytics option.....	86
VxBlock Central architecture.....	87
Data center architecture.....	91
AMP overview.....	93
AMP hardware components (VMware vSphere 6.x).....	94
AMP software components (VMware vSphere 6.x).....	94
Management network connectivity (VMware vSphere 6.x).....	95
Virtualization layer (VMware vSphere 6.x).....	99
Virtualization components (VMware vSphere 6.x).....	99
VMware vSphere Hypervisor ESXi (VMware vSphere 6.x).....	99
Cluster configurations.....	99
Datstores.....	100
Virtual networks.....	100
VMware vCenter Server (VMware vSphere 6.x).....	100
VxBlock System 1000 sample configurations.....	103
VxBlock System 1000 with AMP Central integrated configuration.....	103

VxBlock System 1000 with AMP-3S..... 105

VxBlock System 1000 with Dell EMC Unity 350F storage.....109

VxBlock System 1000 with Dell EMC Unity XT storage..... 110

VxBlock System 1000 with Dell EMC Unity and Isilon Gen 6 storage..... 111

VxBlock System 1000 with PowerMax 2000 storage..... 114

VxBlock System 1000 with PowerMax 8000 storage..... 116

VxBlock System 1000 with VMAX 950F storage..... 119

VxBlock System 1000 with XtremIO X2 storage..... 122

Revision history

Date	Document revision	Description of changes
December 2019	1.14	Updated the document with the following: <ul style="list-style-type: none"> • Added support for AMP Central. • Added support for VxBlock Central release 2.5. Added support for VxBlock Central Life Cycle Management • Added support for PowerMaxOS 5978.444. • Removed 2012 from Windows Server reference.
September 2019	1.13	Added support for Dell EMC Unity XT storage. Added information about VMware vSphere 6.7 and 6.5 with regard to external VMware PSCs. Add note that states as of RCM 6.5.10.0 and 6.7.3.0, Cisco DCNM is no longer supported.
September 2019	1.12	Updated <i>VxBlock Central Workflow Automation option</i> .
July 2019	1.11	Updated for VxBlock Central Version 2.0. <ul style="list-style-type: none"> • SRS configuration feature in the VxBlock Central user interface. • Component discovery feature in the VxBlock Central user interface. • New VxBlock Central alerts.
June 2019	1.10	<ul style="list-style-type: none"> • Added support for expanding an AMP-3S cluster with Cisco UCS C220 M5 servers. • Added support for Cisco UCS VICs 1495 and 1497. • Added support for the Cisco MDS 9148T and 9396T 32G Multilayer Fabric Switches. • Updated for new VxBlock Central licensing.
March 2019	1.9	Added the following features: <ul style="list-style-type: none"> • AMP-3S with Cisco UCS C220 M5 servers and VMware VDS • VxBlock Central features • VMware vSphere 6.7 • Dell EMC Unity Hybrid storage arrays • Cisco UCS C480 M5 Server • Cisco UCS Gen 4
December 2018	1.8	Added support for VMware vSphere Distributed Switch and VxBlock Central.
November 2018	1.7	Updated <i>Cisco UCS Servers</i> for Cisco UCS GPUs.
August 2018	1.6	Added support for PowerMax 2000 and 8000 storage arrays. Added support for AMP-3S. Added support for the Cisco Nexus 9336C-FX2 Switch.

Revision history

Date	Document revision	Description of changes
July 2018	1.5	For Cisco UCS B480 M5 Blade Server: <ul style="list-style-type: none"><li data-bbox="523 331 895 363">• Updated <i>Compute components</i>.<li data-bbox="523 375 863 407">• Updated <i>Cisco UCS Servers</i>.
May 2018	1.4	Updated the section <i>Dell EMC Unity Software licensing</i> with information about data reduction.
May 2018	1.3	Added support for XtremIO X2.
March 2018	1.2	Added support for Cisco UCS C-Series M5 servers.
February 2018	1.1	Updated graphics in the compute section and the sample elevations.
February 2018	1.0	Initial version

Introduction

This document describes the high-level design of the VxBlock System 1000 and the hardware and software components. The VxBlock 1000 offers the latest storage arrays, compute, and networking equipment from Cisco Systems and AMP management infrastructure.

Customization

The VxBlock 1000 enables a choice of multiple storage arrays, compute servers, and network stacks. The compute layer includes both Cisco UCS B-Series and C-Series Servers. The storage layer includes multiarray support from VMAX All-Flash, Dell EMC Unity, Dell EMC Unity XT, PowerMax, XtremIO, and Isilon storage to fulfill both block and file storage needs.

Networking options for VxBlock 1000 with Cisco UCS Gen 3 include:

- LAN networking supports 10 Gb/s, 40 Gb/s, and 100 Gb/s uplinks to the network. LAN networking supports 10 Gb/s or 40 Gb/s access for attached devices, depending on the configuration.
- SAN networking supports 16 Gb/s FC access from both the compute and storage arrays.

Networking options for VxBlock 1000 with Cisco UCS Gen 4 include:

- LAN supports 10 Gb/s, 40 Gb/s, and 100 Gb/s uplinks to the network. LAN supports 10 Gb/s or 25 Gb/s access for attached devices, depending on the configuration.
- SAN supports 16 Gb/s or 32 Gb/s FC access from the compute components. SAN supports 16 Gb/s FC access from the storage arrays.

Expansion

You can expand the VxBlock 1000 by adding compute, network, and storage resources. After the initial deployment, you can add more disks/arrays for increased capacity and servers to increase performance.

Management

VxBlock Central user interface is provided in multiple options to manage VxBlock Systems.

AMP-VX is scalable and deployed inside the VxBlock System 1000. AMP-VX can manage multiple system platforms in a single data center.

AMP Central provides a central/multisystem management platform with Cisco UCS C220 M5 Servers and Dell EMC Unity 300 Hybrid for shared storage. AMP Central is available in stand-alone and integrated configurations.

AMP Central integrated is scalable and deployed inside the VxBlock System or in a stand-alone configuration with dedicated cabinets and dedicated management and ToR switches. AMP Central can manage multiple system platform VMware vCenter Servers in a single logical data center.

AMP-3S is a single system management infrastructure platform to manage a single platform in a data center.

System overview

The VxBlock System 1000 enables a choice of multiple storage arrays, compute servers, and network stacks.

The following table provides key hardware and software components for management, virtualization, compute, network, and storage for VxBlock 1000:

Resource	Components
Management interface	<p>Vision Intelligent Operations or VxBlock Central manages and monitors VxBlock Systems</p> <p>VxBlock Central provides the following options:</p> <ul style="list-style-type: none"> • The VxBlock Central Base option provides the VxBlock Central user interface. • The VxBlock Central Advanced Analytics option provides access to VxBlock Central Operations capabilities to monitor the health status and metrics with VMware vRealize Operations (vROps) Manager.
Management platforms	<p>AMP-VX, AMP Central, or AMP-3S</p> <p>The following components are available for management:</p> <ul style="list-style-type: none"> • License Manager for PowerPath/VE • PowerPath Management Appliance • Cisco UCS Manager • Cisco UCS CIMC • Unisphere for Dell EMC Unity • Unisphere for VMAX • Isilon OneFS • XtremIO XMS • Secure Remote Services • Cisco Data Center Network Manager (DCNM) for SAN • Cisco DCNM for LAN
Virtualization	<ul style="list-style-type: none"> • VMware vSphere Server Enterprise Plus • VMware vSphere ESXi • VMware vCenter Server • VMware vSphere Web Client • VMware vSphere HTML5 Client • VMware Platform Services Controller (PSC) • PowerPath/VE

Resource	Components
Compute	<ul style="list-style-type: none"> • Cisco UCS 5108 Blade Server Chassis • Cisco UCS B-Series M4 and M5 Servers • Cisco UCS C-Series M4 and M5 Servers <p>Cisco UCS Gen 3 components:</p> <ul style="list-style-type: none"> • Cisco UCS 2304 Fabric Extenders • Cisco Nexus 2348UPQ Fabric Extenders • Cisco UCS 6332-16UP Fabric Interconnects <p>Cisco UCS Gen 4 components:</p> <ul style="list-style-type: none"> • Cisco UCS 2204 or 2208 Fabric Extenders • Cisco Nexus 22232PP Fabric Extenders • Cisco UCS 6454 Fabric Interconnects
Network	<ul style="list-style-type: none"> • Cisco Nexus 9336C-FX2 Data Plane Switch • Cisco MDS 9148S/9148T Multilayer Fabric Switch, Cisco MDS 9396S/9396T Multilayer Fabric Switch, Cisco MDS 9710 Multilayer Director, or Cisco MDS 9706 Multilayer Director • Cisco Nexus 31108TC-V Switches • Optional Cisco Nexus 3232C or Cisco Nexus 9336C-FX2 Management Aggregation Switches • VMware vSphere Distributed Switch (VDS) <p>Cisco UCS Gen 3 components:</p> <ul style="list-style-type: none"> • Cisco Nexus 93180YC-FX, Cisco Nexus 93180LC-EX Data Plane Switches
Storage	<ul style="list-style-type: none"> • VMAX All-Flash 950F and 950FX • VMAX All Flash 250F and 250FX • Dell EMC Unity 300, 400, 500, 600, and 350F, 450F, 550F, 650F • Dell EMC Unity XT 380, 480, 680, 880, and 380F, 480F, 680F, 880F • Isilon F800, H600, H500, H400, A200, and A2000 • PowerMax 2000 • PowerMax 8000 • XtremIO X2-R and X2-S

See the *Dell EMC AMP-VX Product Guide* or the *Dell EMC AMP Central Product Guide* for more information about AMP-VX.

Minimum and maximum configurations

VxBlock System 1000 compute, network, and storage components scale with blade and rack mount servers, bandwidth and connectivity options, and storage capacity and performance.

Compute

The following table lists the minimum and maximum compute requirements:

Component	Minimum	Maximum	
		Cisco UCS Gen 3	Cisco UCS Gen 4
Cisco UCS domain	1	7 with Cisco Nexus 9336-FX2 Switch	7
Cisco UCS 5108 Blade Server Chassis	2 (unless using a Cisco UCS C-Series server configuration)	10 per domain (total of 50)	7 across domains (total of 77)
Cisco UCS B-Series or C-Series servers	3	8 half-width per chassis (total of 400) 20 40 GbE Cisco UCS C-Series Rack Mount servers per domain (total of 140) 160 FEX-connected 10 GbE Cisco UCS C-Series Rack Mount servers per domain (total of 1120)	8 half-width per chassis (total of 616) 44 w/ dual port, 22 total with quad port 25 GbE Cisco UCS C-Series M5 only Rack Mount servers per domain (total of 140) 160 FEX-connected 10 GbE Cisco UCS C-Series Rack Mount servers per domain (total of 1120)
Cisco UCS C-Series server for bare metal deployments	1	N/A	N/A

Cisco UCS B-Series blade type including all supported M4 and M5 blade configurations. Cisco UCS C-Series servers including all supported M4 and M5 configurations.

Network

The following table lists the minimum and maximum network requirements.

Cisco Nexus Switch	Minimum	Maximum
Cisco Nexus 31108TC-V Switches	1 pair (96 management ports)	2 pairs (192 management ports)
Cisco Nexus 3232C or Cisco Nexus 9336C-FX Management Aggregation Switches	N/A	1 pair
Cisco MDS 9148S/9148T Multilayer Fabric Switches	1 pair	N/A
Cisco Nexus 9336C-FX data plane switch	1 pair	10 pairs
Cisco MDS 9710 Multilayer Directors with 348 16 Gb/s or 32 Gb/s FC ports each	N/A	1 pair

Storage: VMAX All-Flash

The following table lists the minimum and maximum storage requirements for VMAX All-Flash

VMAX All-Flash storage	Component	Minimum	Maximum
950F and 950FX	Engine	1	8
	Front-end ports	16	192
	Drive support	17	1920
250F and 250FX	Engine	1	2
	Front-end ports	16	64
	Drive support	9	100

Note: A single-cabinet configuration of the VMAX All-Flash 250F and 250FX is available with a single engine.

Storage: PowerMax

The following table lists the minimum and maximum storage requirements for PowerMax.

PowerMax storage	Component	Minimum	Maximum
8000	Engine	1	8
	Front-end ports	16	256
	Drive support	8	288
2000	Engine	1	2
	Front-end ports	16	64
	Drive support	5	96

Note: A single-engine configuration of the PowerMax 8000 is limited to 3 IOMs per director.

Dell EMC Unity and Dell EMC Unity XT

The following table provides minimum and maximum drives for Dell EMC Unity and Dell EMC Unity XT

Dell EMC Unity storage array	Minimum drives	Maximum drives
350F All-Flash array	10	150
450F All-Flash array	10	250
550F All-Flash array	10	500
650F All-Flash array	10	1000
XT 380F All-Flash array	10	500
XT 480F All-Flash array	10	750
XT 680F All-Flash array	10	1000
XT 880F All-Flash array	10	1500
300 Hybrid arrays	10	150
400 Hybrid arrays	10	250
500 Hybrid arrays	10	500
600 Hybrid arrays	10	1000
380F XT Hybrid arrays	10	500
480F XT Hybrid arrays	10	750
680F XT Hybrid arrays	10	1000
880F XT Hybrid arrays	10	1500

Note: The storage array count depends on the number of 16 Gb/s FC ports available in a single Cisco MDS director per SAN fabric.

XtremIO X2

The following table lists the minimum and maximum storage requirements for X2.

XtremIO storage	Component	Minimum	Maximum
X2	Drives per X-Brick	18	72
	Drives per cluster	72	288
X2-S X-Brick	X-Brick	7.2 TB minimum raw	28.8 TB maximum raw Up to 4 X-Brick blocks per cluster
	FC ports (16 Gb/s)	4	16
	Ethernet ports (10 Gb/s) for management	2	2
	Active/active controllers	2	8
X2-R X-Brick	X-Brick minimum raw	34.5 TB minimum raw	230 TB maximum raw Up to 4 X-Brick blocks per cluster
	FC ports (16 Gb/s)	4	16
	Ethernet ports (10 Gb/s) for management	2	2
	Active/active controllers	2	8

Storage: Isilon

The following table provides minimum and maximum storage drives for Isilon.

Model	Minimum drives	Maximum drives
F800 and H600	2.5 inches	2.5 inches
H500, H400, A200, and A2000	4 inches	3.5 inches

AMPs

You can configure AMP Central with four to 16 servers to manage up to eight production VMware vCenter Servers. AMP Central has a dedicated VMware vCenter Server for the local management workload and dedicated VMware vCenter Servers for the production workload. An integrated AMP Central supports up to four production VMware vCenter Servers, while AMP Central stand-alone supports up to eight production VMware vCenter Servers. Size AMP Central servers and Dell EMC Unity storage to support VMware vCenter Servers and data protection requirements. AMP Central supports core, optional, and ECO management workloads.

You can configure AMP-3S with two to six Cisco UCS M4 servers or 4 to 16 Cisco UCS M5 servers. Three servers are required to build a AMP-3S M4 cluster with VMware vSphere HA and DRS cluster. The AMP-3S M4 cluster is based on the core and optional workloads. AMP-3S M5 supports the additional ECO management workload. AMP-3S manages a single VxBlock System.

AMP-VX manages up to eight VxBlock Systems. AMP-VX contains a single instance of VxBlock Central, VxBlock Central Orchestration Services, and VMware vRealize Orchestrator. The number of VxBlock Systems that VxBlock Central Orchestration consumes does not impact the number of VxBlock Central instances. Up to five servers can be deployed in a single workflow.

The combined components offer balanced CPU, I/O bandwidth, and storage capacity relative to the compute and storage arrays in the VxBlock 1000. All components have N+N or N+1 redundancy.

The following table shows the maximum number of Cisco UCS servers with seven domains and connected ToR switches for Cisco UCS FIs.

Fabric interconnects	B-Series half width	C-Series direct-connected	C-Series FEX-connected
Cisco UCS 6332-16UP	400	140	1120

Fabric interconnects	B-Series half width	C-Series direct-connected	C-Series FEX-connected
Cisco UCS 6454	616	308	1120

Scale up compute resources

To scale up compute resources, add servers and chassis when the VxBlock System 1000 is built or after it is deployed.

Servers

Cisco UCS servers are sold individually. Each server type must have a minimum of three servers as a base and can increase in single server increments. Each server is added with license packs for the following software:

- Cisco UCS Manager
- VMware vSphere ESXi
- PowerPath/VE

The base configuration of a VxBlock 1000 includes a minimum of three servers. The maximum number of servers depends on the selected scale point.

License packs for VMware vSphere ESXi and PowerPath/VE are not available for bare metal blades.

Chassis

The power supplies and FEX for all chassis are prepopulated and cabled, and all required Twinax cables and transceivers are populated. However, in the base configuration, there is a minimum of two Cisco UCS 5108 Blade Server Chassis. There is no unpopulated server chassis unless they are ordered that way. This limited licensing reduces the entry cost for the VxBlock 1000.

As more blades are added and more chassis are required, more chassis are added automatically to an order. The kit contains software licenses to enable extra FI ports.

Only enough port licenses for the minimum number of chassis to contain the blades are ordered. More chassis can be added to enable flexibility in the field or to initially spread the blades across a larger number of chassis.

Cisco Nexus 2300 Fabric Extenders (Cisco UCS Gen 3 configurations)

For Cisco Nexus FEX-connected C-Series servers, the first pair of FEXs can support up to 48 servers. Each additional pair of FEXs can support another 48 servers until the fourth pair of FEXs has been added. The fourth pair, FEX 7, and FEX 8, can support only 16 FEX-connected Cisco UCS C-Series servers. With eight FEXs connected, the Cisco UCS domain can support the Cisco UCS Manager maximum of 160 servers.

Each individual FEX requires two uplinks to the FI for every 16 connected servers. Four uplinks per FEX are required to support 17–32 servers, and six uplinks are required to support 33–48 Cisco UCS C-Series servers.

Populate each pair of Cisco Nexus FEXs with 48 connected servers before adding a new pair to the domain.

Cisco Nexus 2200 Fabric Extenders (Cisco UCS Gen 4 configurations)

For Cisco Nexus FEX-connected C-Series servers, the first pair of FEXs can support up to 32 servers. Each additional pair of FEXs can support another 32 servers for up to five pairs. With 10 FEXs connected, the Cisco UCS domain can support the Cisco UCS Manager maximum of 160 servers.

Each FEX requires four or eight uplinks to the fabric interconnect. With 4 Uplinks the FEX can support 8–16 connected servers. Eight uplinks are required per FEX to support 24–32 Cisco UCS C-Series servers.

Populate each pair of Cisco Nexus FEXs with 32 connected servers before adding a new pair to the domain.

Network topology

LAN and SAN connectivity is segregated into separate Cisco switches for network connectivity.

The following table shows the switches for LAN and SAN connectivity:

LAN connectivity (Cisco UCS Gen 3)	LAN connectivity (Cisco UCS Gen 4)	SAN Connectivity
<ul style="list-style-type: none"> • Cisco Nexus 93180YC-FX Switch • Cisco Nexus 93180LC-EX Switch • Cisco Nexus 9336C-FX2 Switch 	<p>Cisco Nexus 9336C-FX2 switch</p>	<ul style="list-style-type: none"> • Cisco MDS 9148S 16 Gb/s Multilayer Fabric Switch • Cisco MDS 9396S 16 Gb/s Multilayer Fabric Switch • Cisco MDS 9148T 32 Gb/s Multilayer Fabric Switch • Cisco MDS 9396T 32 Gb/s Multilayer Fabric Switch • Cisco MDS 9706/9710 Multilayer Director with 48 port 16 Gb/s line cards and/or 48 port 32 Gb/s line cards

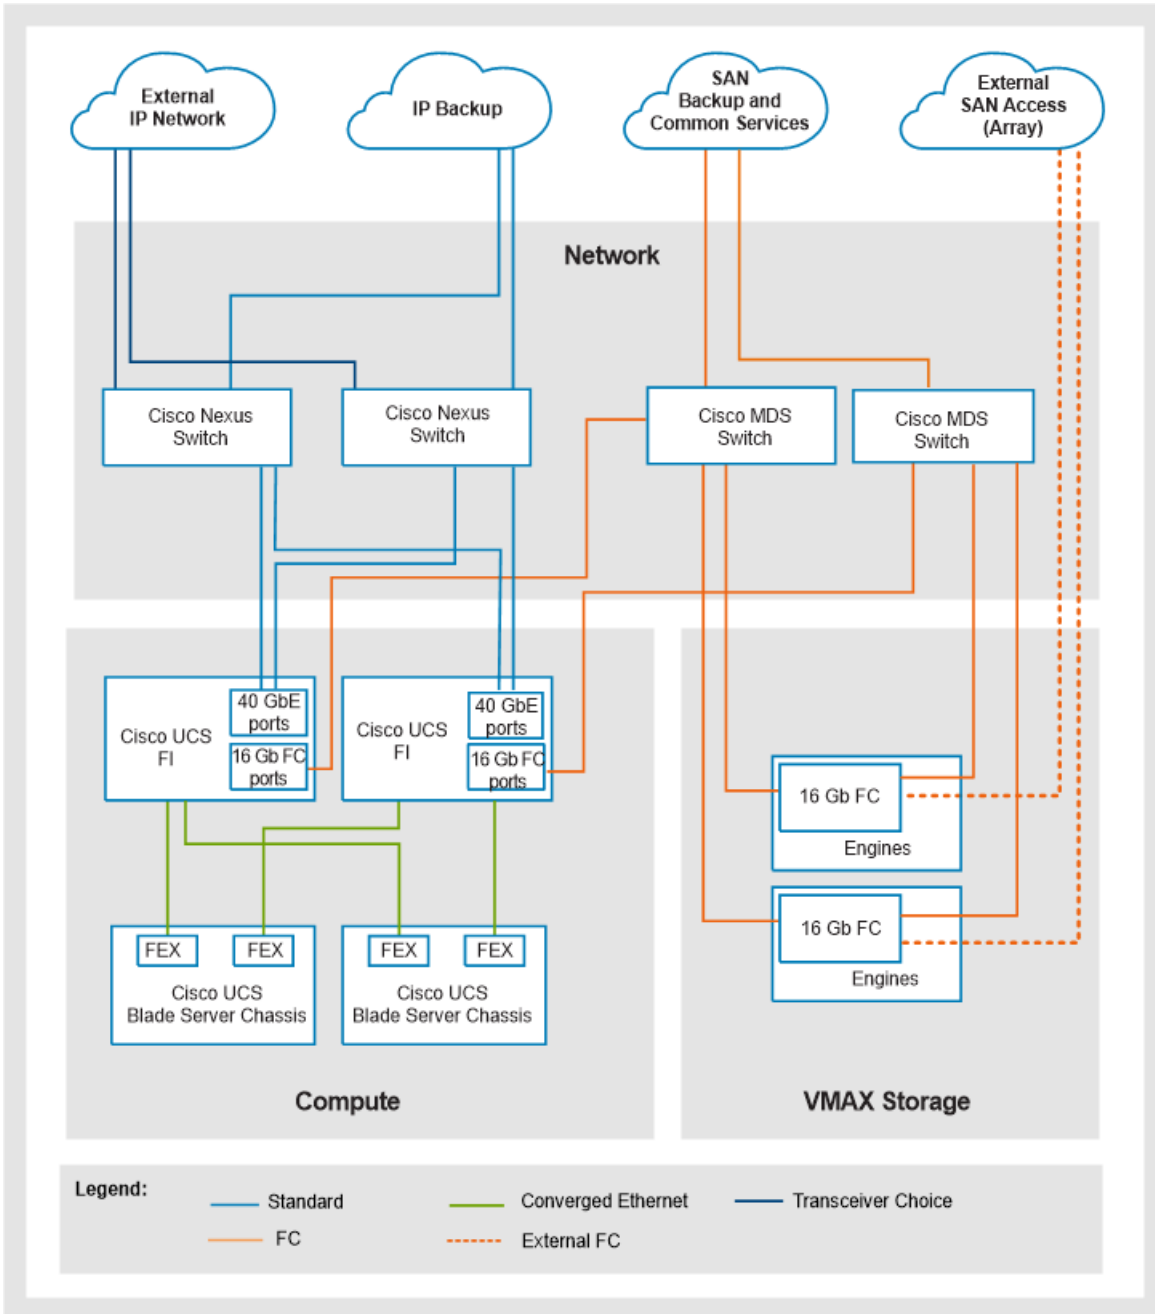
The compute layer connects to Ethernet and FC components in the network layer. Cisco UCS FIs connect to the Cisco Nexus switches in the Ethernet network through 40 GbE or 10 GbE port channels.

Cisco UCS Gen 3 FIs connect to the Cisco MDS switches through port channels that are made up of multiple 16 Gb/s links.

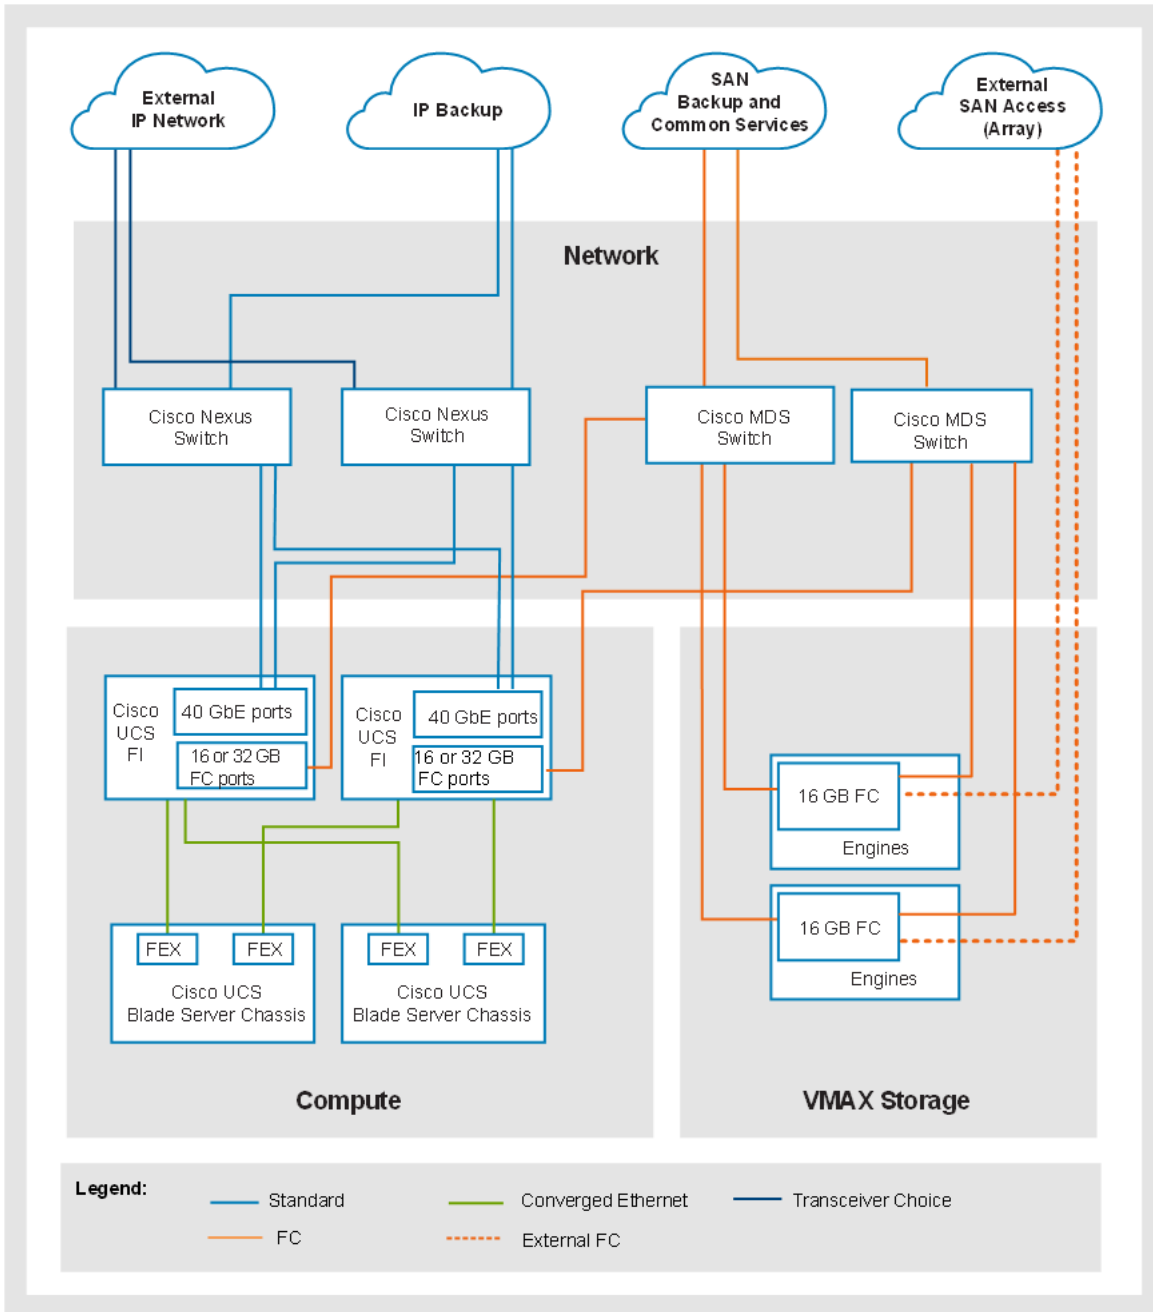
Cisco UCS Gen4 FIs connect to the Cisco MDS switches through port channels that are made up of multiple 16 Gb/s or 32 Gb/s links. The Cisco MDS switches autonegotiate the port speed based on the SFPs in the Cisco UCS FIs.

Front-end I/O modules in the storage array connect to the Cisco MDS switches in the network layer over 16 Gb/s FC links.

The following figure shows a network block storage configuration for the VxBlock System 1000:



The following figure shows a network block storage configuration for the VxBlock System 1000 with 32 Gb/s SAN switches:



Host boot configuration

By default, VMware vSphere ESXi boots from SD cards.

VMware vSphere ESXi hosts can optionally boot over the FC SAN.

Compute layer

Cisco UCS Servers provide computing power within the VxBlock System 1000.

Component	Description
Cisco UCS 6332-16UP Fabric Interconnects	<ul style="list-style-type: none"> Contains four or six 40 Gb/s links Connects to a pair of 40 Gb/s capable Cisco Nexus 9336C-FX2 or 93180LC-EX Switches Connects to a pair of 10 Gb/s capable Cisco Nexus 93180YC-FX Switches using 4 x 10 Gb/s break-out cables Configured with eight, 12, or 16 FC connections per FI
Cisco UCS Gen 4 6454 Fabric Interconnects	<ul style="list-style-type: none"> Contains two, four, or six 100 Gb/s links Connects to a pair of 100 Gb/s capable Cisco Nexus 9336C-FX2 Switches Configured with four or eight FC connections per FI

Each FI has multiple ports that are reserved for FC ports. These ports connect to Cisco SAN switches and carry FC traffic between the compute layer and the storage layer. SAN port channels carry FC traffic between the FIs and upstream Cisco MDS switches.

FEX within the Cisco UCS 5108 Blade Server Chassis connect to FIs over converged networking.

Cisco UCS Gen 3	Cisco UCS Gen 4
<p>Two or four 40 GbE connections on each FEX connect northbound to the FIs, regardless of the number of blades in the chassis.</p> <p>Cisco UCS C-Series Servers can connect directly to ports on Cisco UCS 6332-16UP FIs.</p> <p>Optionally, a pair of Cisco Nexus 2348UPQ FEXs can be used for increased Cisco UCS C-Series server scalability. These connections carry IP and FC traffic.</p>	<p>Four or eight 10 GbE connections on each FEX connect northbound to the FIs, regardless of the number of blades in the chassis.</p> <p>Cisco UCS C-Series servers can connect directly to ports on the Cisco UCS 6454 FIs.</p>

By default, servers in the VxBlock 1000 boot VMware vSphere ESXi from local mirrored SD cards. Booting from SAN is optional.

Component	Model
Cisco UCS Gen 3 FIs	Cisco UCS 6332-16UP Fabric Interconnect (required)
Cisco UCS Gen 4 FIs	Cisco UCS 6454 Fabric Interconnect (required)
Cisco UCS Gen 3 FEX	Cisco UCS 2304 Fabric Extender, Cisco Nexus 2348UPQ Fabric Extender
Cisco UCS Gen 4 FEX	Cisco UCS 2204/2208 Fabric Extender, Cisco Nexus 2232PP Fabric Extender
Chassis	Cisco UCS 5108 Chassis
B-Series servers	Cisco B200 M4, B260 M4, B420 M4, B460 M4, B200 M5, B480 M5 Servers
C-Series servers	Cisco C220 M4, C240 M4, C220 M5, C240 M5, C480 M5 Servers

Cisco UCS fabric interconnects

The default Cisco UCS 6332-16UP Fabric Interconnects are 40 port, 16 Gb/s FC, and 40 GbE capable switches. The Cisco UCS FIs run the Cisco UCS Manager application and database. Each VxBlock System 1000 contains at least two FIs that form a high availability cluster.

Cisco UCS 6332-16UP Fabric Interconnects provide network connectivity and management capability to the Cisco UCS blades and chassis. Each pair of Cisco UCS 6332-16UP Fabric Interconnects form a Cisco UCS compute domain.

The following table shows how many Cisco UCS compute domains the Cisco Nexus ToR switches can support:

Component	Maximum Cisco UCS compute domains
Cisco Nexus 93180YC-FX	2
Cisco Nexus 93180LC-EX	5
Cisco Nexus 9336C-FX2	7

See to the *Cisco UCS 6300 Series Fabric Interconnect Specification Sheet* for additional information regarding Cisco UCS 6332-16UP Fabric Interconnects.

The Cisco UCS 6454 Gen 4 Fabric Interconnects are 48 port, 16 Gb/s, or 32 Gb/s FC, 10/25 GbE and 6 QSFP 100 GbE capable switches that run the Cisco UCS Manager application and database. Each VxBlock 1000 contains at least two FIs, which form a high availability cluster.

Cisco UCS Gen 4 Cisco UCS 6454 Fabric Interconnects provide network connectivity and management capability to the Cisco UCS blades and chassis. Each pair of Cisco UCS 6454 FIs form a Cisco UCS compute domain. The Cisco Nexus 9336C-FX2 Switches can support a maximum of seven Cisco UCS compute domains.

Cisco UCS servers

The Cisco UCS data center platform unites compute, network, and storage access.

Optimized for virtualization, the Cisco UCS Gen 3 integrates a low-latency, lossless 40 Gb/s or 10 GbE on Cisco UCS 6332-16UP Fabric Interconnects.

Cisco UCS Gen 4 integrates a 10 Gb/s blade server and 10/25 GbE rack server unified network fabric, with enterprise-class, x86-based Cisco UCS B-Series and C-Series servers.

The VxBlock System 1000, powered by Cisco UCS, offers the following features:

- Built-in redundancy for high availability
- Hot-swappable components for serviceability, upgrade, or expansion
- Fewer physical components than in a comparable system built piece by piece
- Reduced cabling
- Improved energy efficiency over traditional blade server chassis

For VMware deployments, the VxBlock System 1000 can scale from a minimum of three Cisco UCS B-Series blade servers or three Cisco UCS C-Series Servers, to hundreds of Cisco UCS B-Series or C-Series servers. The maximum number of servers that can be installed depends on the number of Cisco UCS domains in the system and the underlying storage array configuration.

The maximums assume that:

- Server models are consistent within the domain/system.
- The configured host count is supported by the connected storage arrays.

Mixing server models in a domain or system is supported but reduces the maximum quantity. Cisco UCS B-series blade servers require 10 Cisco UCS chassis per domain with the two-link connection option to achieve the maximum. The maximum quantity that supported per system can only be achieved with five Cisco UCS domains.

The following table lists the maximum supported Cisco UCS servers in a Cisco UCS Gen 3 environment:

Server model and type	Maximum servers per Cisco UCS domain	Maximum servers per system w/ Cisco Nexus 93180 Switch
Cisco UCS B200 M4, B200 M5 (half-width)	80	400
Cisco UCS B420 M4, B260 M4, B480 M5 (full-width)	40	200
Cisco UCS B460 M4 (double-height)	20	100
Cisco UCS C220 M4, C240 M4, C220 M5, C240 M5, C480 M5 (Direct connect - 40 Gb/s)	20	100
Cisco UCS C220 M4, C240 M4, C220 M5, C240 M5, C480 M5 (FEX connect - 10 Gb/s)	160	800

The following table lists the maximum supported Cisco UCS servers in a Cisco UCS Gen 4 environment:

Server model and type	Maximum servers per Cisco UCS domain	Maximum servers per system w/ Cisco Nexus 9336 Switch
Cisco UCS B200 M4, B200 M5 (half-width)	88	616
Cisco UCS B420 M4, B260 M4, B480 M5 (full-width)	44	308
Cisco UCS B460 M4 (double-height)	22	154
Cisco UCS C220 M4, C240 M4, C220 M5, C240 M5, C480 M5 (Direct connect - 25 Gb/s)	44	308
Cisco UCS C220 M4, C240 M4, C220 M5, C240 M5 (FEX connect - 10 Gb/s)	160	1120

The following table shows the maximum Cisco UCS B-Series blade quantities to deploy Cisco UCS C-Series direct-connect servers within the same Cisco UCS Gen 3 domain:

Cisco UCS chassis count	Maximum Cisco UCS B-Series servers	Maximum Cisco UCS C-Series direct-connect servers
0	0	20
1	8	18
2	16	16
3	24	14
4	32	12
5	40	10
6	48	8
7	56	6
8	64	4
9	72	2

Cisco UCS chassis count	Maximum Cisco UCS B-Series servers	Maximum Cisco UCS C-Series direct-connect servers
10	80	0

The following table shows the maximum Cisco UCS B-Series blade quantities for deploying Cisco UCS C-Series direct-connect servers within the same Cisco UCS Gen 4 domain:

Cisco UCS chassis count	Maximum Cisco UCS B-Series servers	Maximum Cisco UCS C-Series direct-connect servers
0	0	44
1	8	36
2	16	32
3	24	28
4	32	24
5	40	20
6	48	16
7	56	12
8	64	8
9	72	4
10	80	2
11	88	0

The following table shows the maximum Cisco UCS B-Series blade quantities for deploying Cisco UCS C-Series FEX-connect servers within the same Cisco UCS Gen 3 domain:

Cisco UCS chassis count	Maximum Cisco UCS B-Series servers	Maximum Cisco UCS C-Series FEX-connect servers
0	0	160
1	8	144
2	16	128
3	24	112
4	32	96
5	40	80
6	48	64
7	56	48
8	64	32
9	72	16
10	80	0

The following table shows the maximum Cisco UCS B-Series blade quantities for deploying Cisco UCS C-Series FEX-connect servers within the same Cisco UCS Gen 4 domain:

Cisco UCS chassis count	Maximum Cisco UCS B-Series servers	Maximum Cisco UCS C-Series FEX-connect servers
0	0	160
1	8	152
2	16	144
3	24	128
4	32	112
5	40	96
6	48	80
7	56	64
8	64	48
9	72	32
10	80	16
11	88	0

The following table shows the CPUs and memory each Cisco UCS server can support:

Cisco UCS server	Maximum supported CPUs and memory
Cisco UCS B200 M5, C220 M5, and C240 M5 servers	Two Intel Xeon scalable CPUs and 3 TB of memory
Cisco UCS B200 M4, B420 M4, C220 M4, and C240 M4 servers	Two Intel Xeon v4 CPUs and 1.5 TB of memory.
Cisco UCS B260 M4 server	Two Intel Xeon v4 CPUs and 3 TB of memory
Cisco UCS B460 M4, B480 M5, and C480 M5 servers	Four Intel Xeon v4 CPUs and 6 TB of memory

GPUs are available in all supported Cisco UCS B-Series and C-Series M4 and M5 servers, either from the factory or by ordering an additional server through the expansion services for VxBlock Systems.

Dell EMC contacts Cisco support for issues and concerns with GPUs ordered as part of a Converged System.

By default, all Cisco UCS B-Series and C-Series servers are configured with local SD cards for OS boot.

Compute server connectivity

Each Cisco UCS B-Series Blade Server contains at least one physical VIC that passes converged FC and IP network traffic through the chassis mid-plane to the FEXs.

Cisco UCS C-Series rack mount servers contain at least one physical VIC as well. The Cisco UCS C-Series rack mount servers can connect directly to each FI or to a pair of Cisco Nexus FEXs which are connected to the FIs.

Cisco UCS B-Series server connectivity

Half-width M4 blade servers can contain a VIC 1340. Half-width M5 servers can contain a VIC 1440 installed in the motherboard (mLOM) mezzanine slot to connect at a potential bandwidth of 20 Gb/s to each fabric.

Optionally, the following VICs can be installed in the PCIe mezzanine slot next to a VIC 1340 or 1440 to direct non-management network traffic to a separate physical adapter:

- VIC 1380 on M4 and M5 servers

- VIC 1480 on M5 servers

In a Cisco UCS B200 Series Server, both the VIC 1340 and 1380 or 1440 and 1480 can connect at 20 Gb/s to each fabric. Also, the VIC 1340 or 1440 can be installed with a port expander card to achieve native 40 Gb/s connectivity to each fabric. However, the VIC 1440 supports a port expander only with a 2304XP IOM. The VIC 1440 does not support port expanders with Cisco UCS 2204XP and 2208XP IOMs.

Full-width M4 blade servers can be configured to contain a VIC 1340. Optionally, a full-width blade M4 server can be configured with a VIC 1340 and 1380.

Full-width M5 servers can contain a VIC 1440 that can connect at 20 Gb/s to each fabric. A full-width M5 server can be configured with a VIC 1440 and 1480.

The VIC 1340 and 1440 can connect at 20 Gb/s. The VIC 1380 and 1480 can communicate at a maximum bandwidth of 40 Gb/s to each fabric.

Another option is to configure the full-width blade server to contain a VIC 1340 or 1440, a port expander card, and a VIC 1380 or 1480 card. With this configuration, the server's network interfaces each communicate at a maximum bandwidth of 40 Gb/s per fabric.

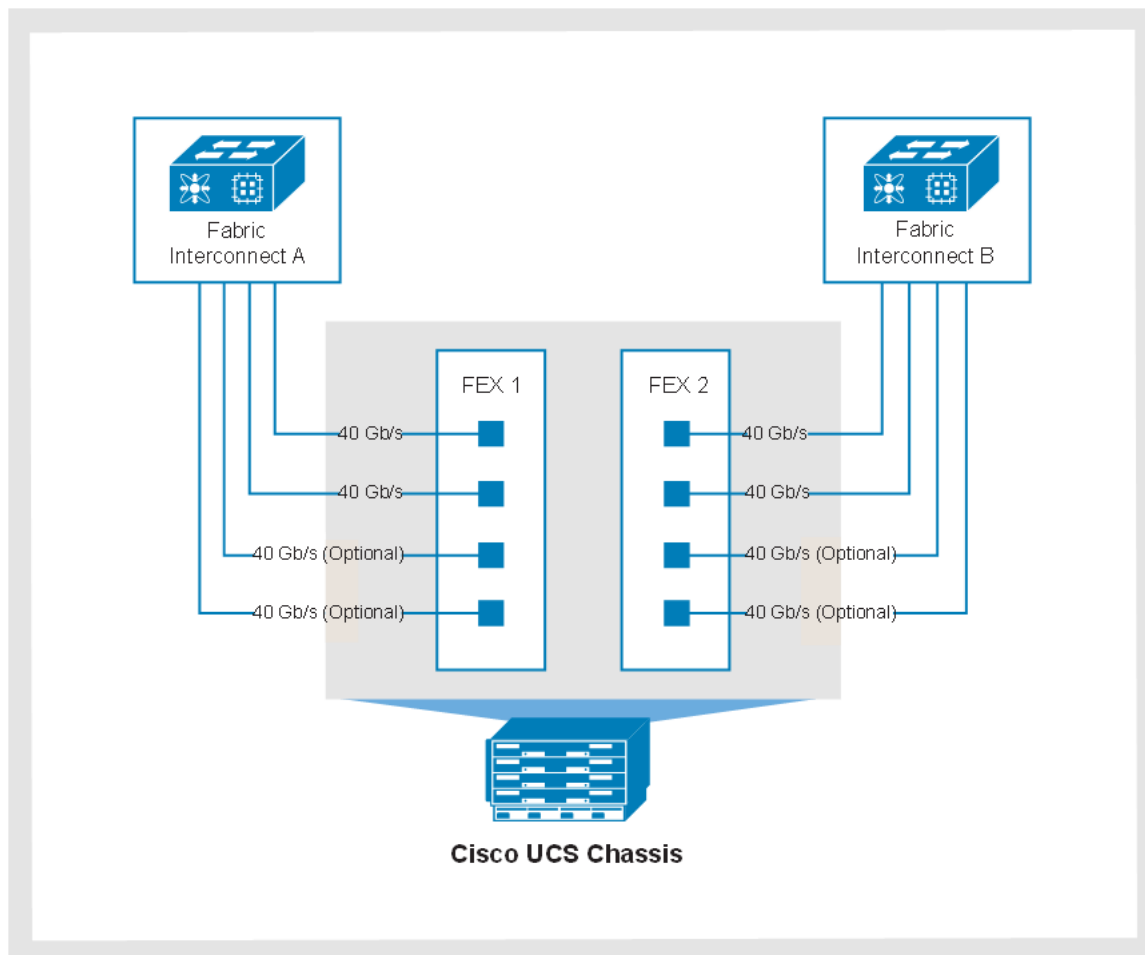
Each Cisco UCS B-Series Server connects to the Cisco UCS 2304 Fabric Extender IO Module (IOM), through the Cisco UCS 5108 Blade Server Chassis midplane. Each IOM carries converged network traffic through two or four 40 Gb/s links to the fabric interconnect.

The following table shows the network connectivity, based on server type and VIC combination:

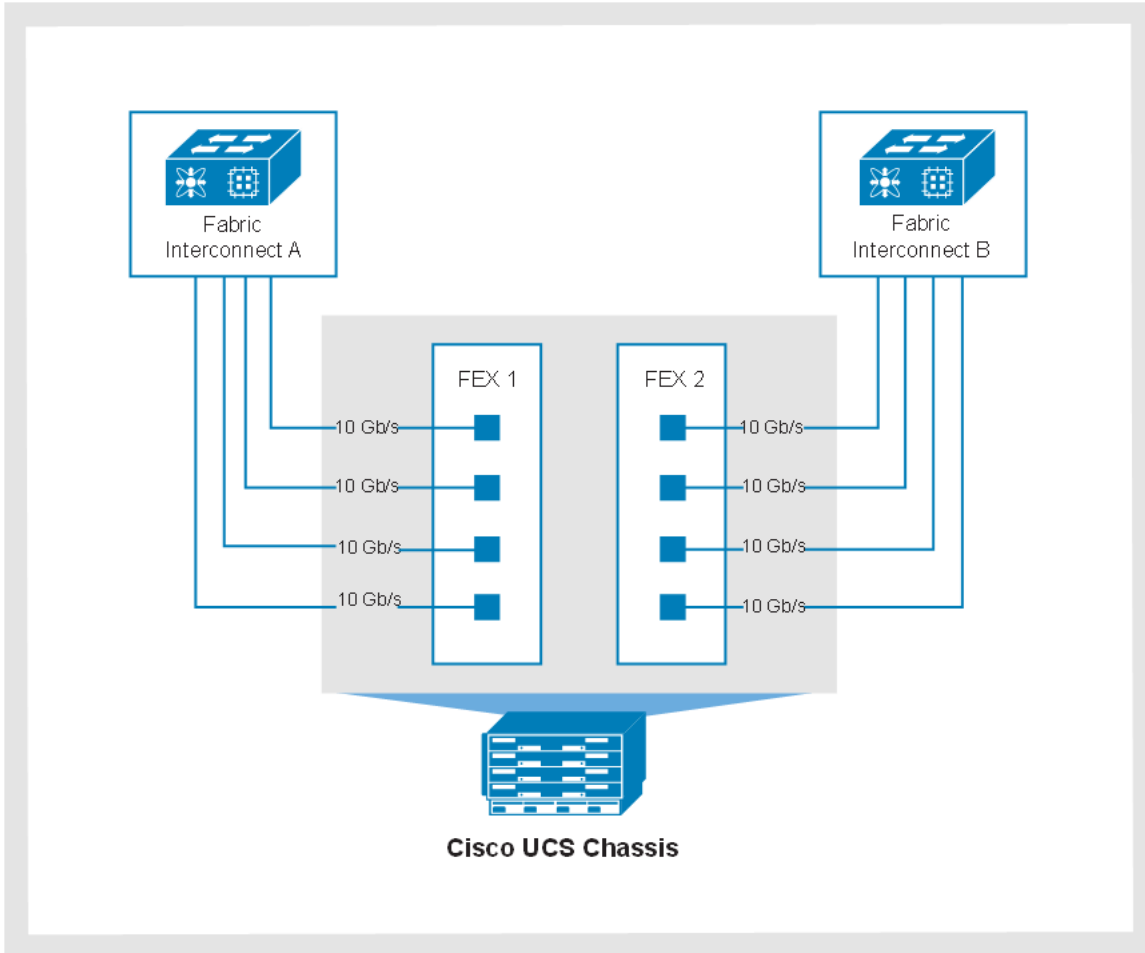
Server type	VIC combination	Bandwidth to each FEX	Total bandwidth available
Half-width blade	1340 only	20 Gb/s	40 Gb/s
	1340 + Port Expander	40 Gb/s	80 Gb/s
	1340 + 1380	20 Gb/s (1340), 40 Gb/s (1380)	80 Gb/s
Full-width blade	1340 only	20 Gb/s	40 Gb/s
	1340 + Port Expander	40 Gb/s	80 Gb/s
	1340 + 1380	20 Gb/s (1340), 40 Gb/s (1380)	120 Gb/s
	1340 + Port Expander + 1380	40 Gb/s (1340), 40 Gb/s (1380)	160 Gb/s
Half-width blade	1440 only	20 Gb/s	40 Gb/s
	1440 + Port Expander	40 Gb/s	80 Gb/s
	1440 + 1480	20 Gb/s (1440), 40 Gb/s (1480)	80 Gb/s
Full-width blade	1440 + 1480	20 Gb/s (1440), 40 Gb/s (1480)	120 Gb/s
	1440 + Port Expander + 1480	40 Gb/s (1440), 40 Gb/s (1480)	160 Gb/s

Note: The VIC 1440 supports a port expander only with a 2304XP. The VIC 1440 does not support port expanders with Cisco UCS 2204XP and 2208XP IOMs.

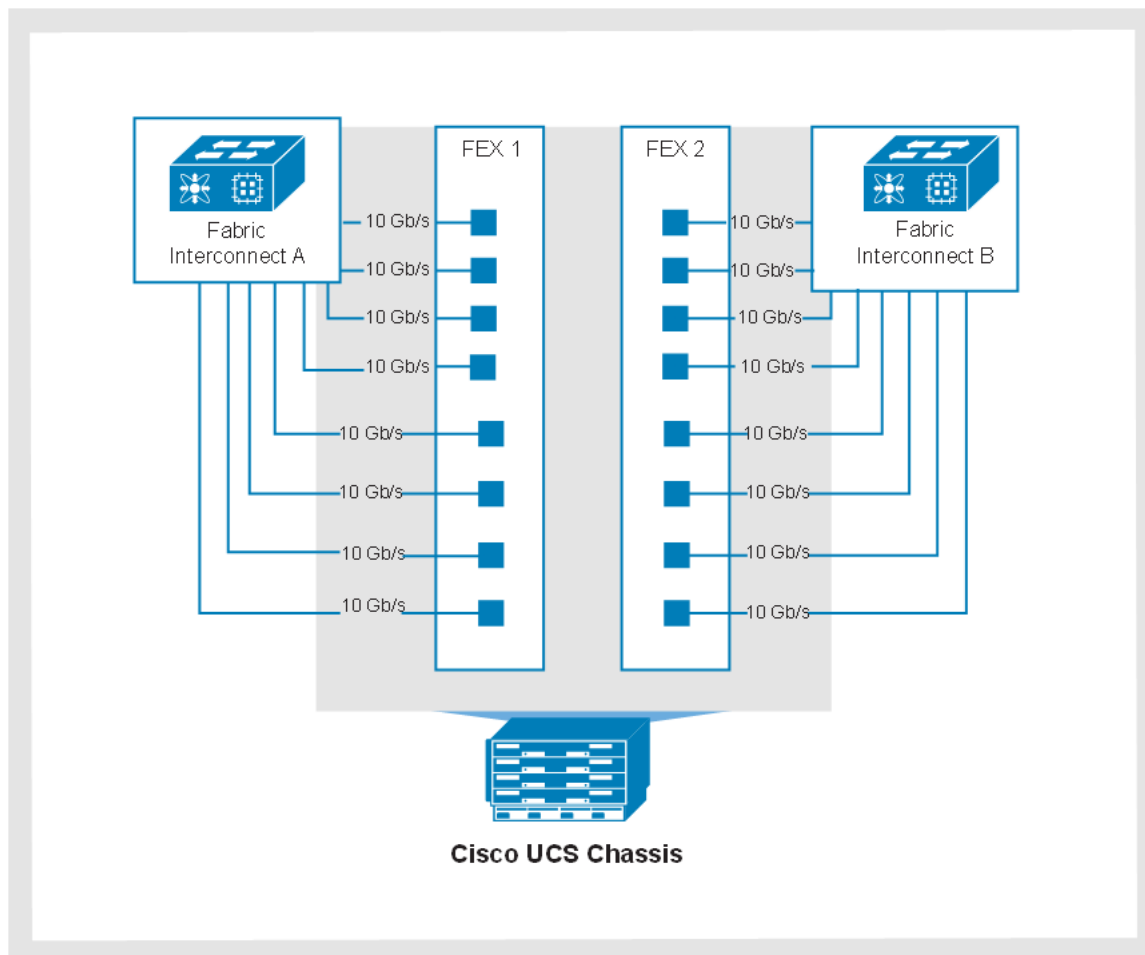
The following figure shows the IO module to Cisco UCS Gen 3 FI connections on a chassis:



The following figure shows a four-link IO module to Cisco UCS Gen 4 FI connections on a chassis:



The following figure shows an eight-link IO module at 10 Gb/s port speed UCS Gen 4 IOM (2208) to Cisco UCS Gen 4 FI (6454) connections on a chassis:



Cisco UCS C-Series Server direct connect

Cisco UCS Manager supports an option to connect the Cisco UCS C-Series Rack Mount server directly to the FIs.

In a Cisco UCS Gen 3 or Gen 4 configuration, you do not need the Cisco Nexus 2300 (Gen 3) or 2200 (Gen 4) Fabric Extenders. This option enables Cisco UCS Manager to manage the Cisco UCS C-Series Rack Mount servers using a single cable for both management and data traffic. The Cisco UCS VIC connects to the FI of the system. This connection uses a single connection from each VIC to each FI. This method is used for configurations that require a few servers.

Note: Different generation VIC adapters are not supported in a single server. For example, in a server with two VIC cards, both VIC adapters should be either 13xx series or 14xx series.

A Cisco UCS Gen 3 C-Series directly connected configuration provides the following features:

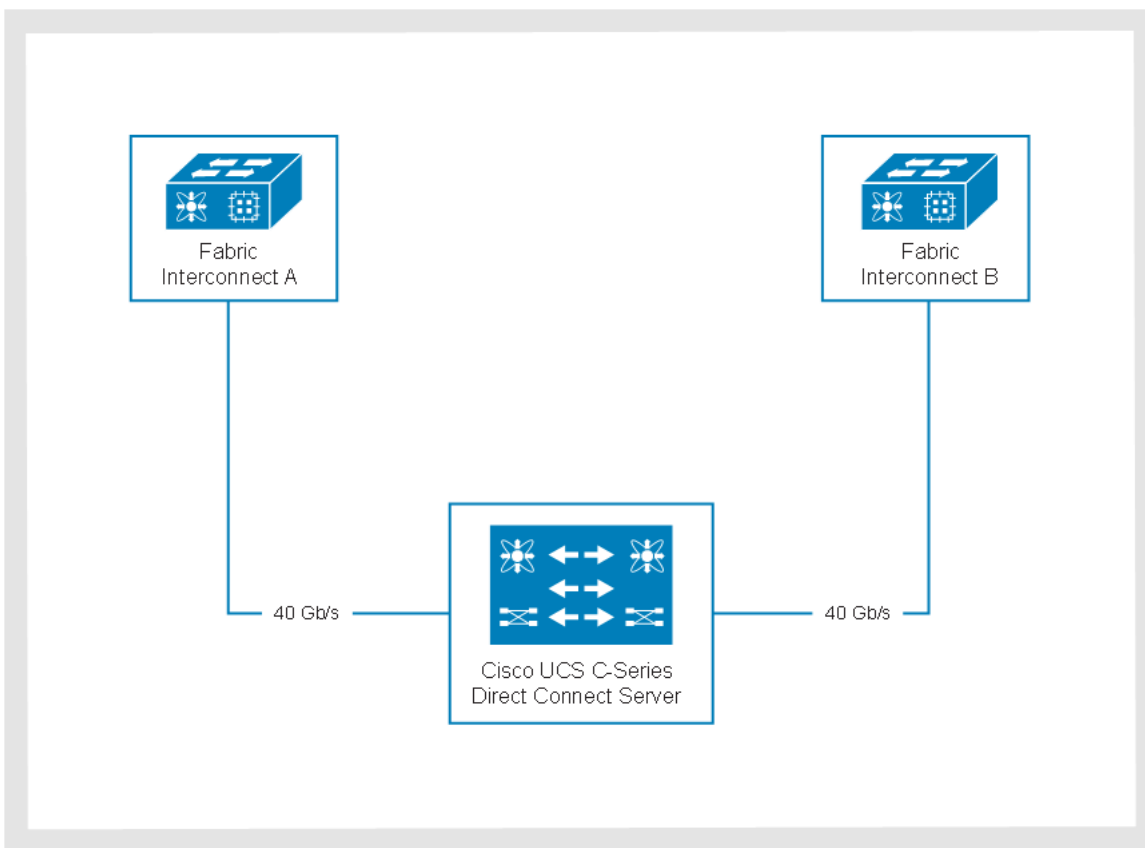
- Cisco UCS C-Series Servers
- 40 Gb/s per fabric from the Cisco UCS VIC 1385 or 1387 on M4 or M5 servers
- 40 Gb/s per fabric from the Cisco UCS VIC 1495 or 1497 on M5 servers only
- Cisco UCS C-Series Server direct connect license

A Cisco UCS Gen 4 C-Series directly connected configuration provides the following features:

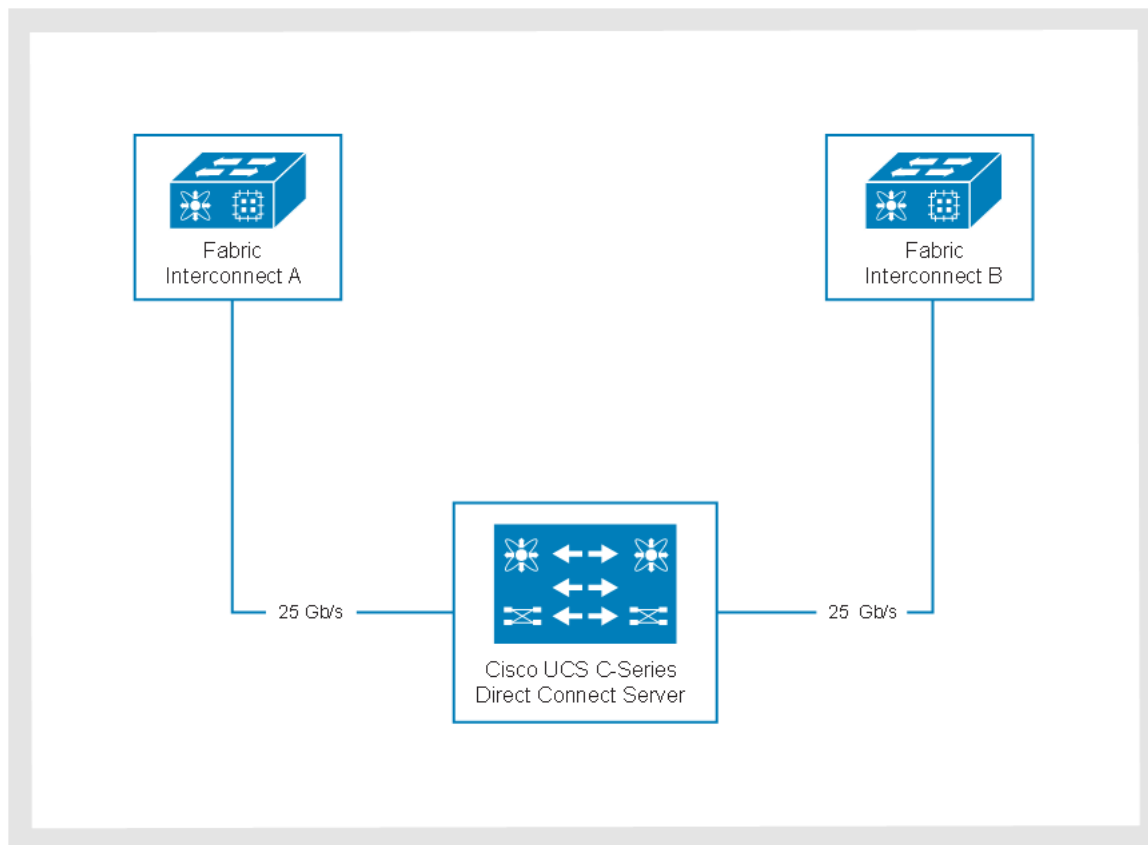
- Cisco UCS C-Series Servers
- 10 Gb/s per fabric from the Cisco UCS VIC 1225 or 1227 on M4 servers and 25 Gb/s VIC 1457 or 1455 on M5 servers

- For the Cisco UCS 1455 and 1457 VICs, ports 1 and 2 are in one port group and connect to Fabric Interconnect A. Ports 3 and 4 are in the other port group and connect to the Fabric Interconnect B.
 - **Cisco UCS 1455/1457 single-link connectivity:** For a two-port connection, port 1 or 2 connects to Fabric Interconnect A and port 3 or 4 connects to Fabric Interconnect B.
 - **Cisco UCS 1455/1457 dual-link connectivity:** For a four-port connection, ports 1 and 2 connect to Fabric Interconnect A and ports 3 and 4 connect to Fabric Interconnect B.
- Note:** Dual-link connectivity is supported on Cisco UCS Gen 4 Fabric Interconnects. However, the Cisco UCS Gen 3 Fabric Interconnect supports only single-link port channels for FC/FCoE between VIC 1455 or 1457 adapters that are on direct-connect rack-servers.
- Cisco UCS C-Series Server direct connect license

The following figure shows a direct connection between the Cisco UCS C-Series server and the Cisco UCS 6332-16UP Fabric Interconnects:



The following figure shows a direct connection between the Cisco UCS C-Series server and the Cisco UCS 6454 Fabric Interconnects:



When a Cisco UCS rack server with Cisco UCS VIC 1455 is connected to a Cisco UCS 6454 FI through a 25 Gb/s cable, manually configure the server port. Port autodiscovery is not supported.

Cisco UCS C-Series Gen 3 Fabric Extender connection

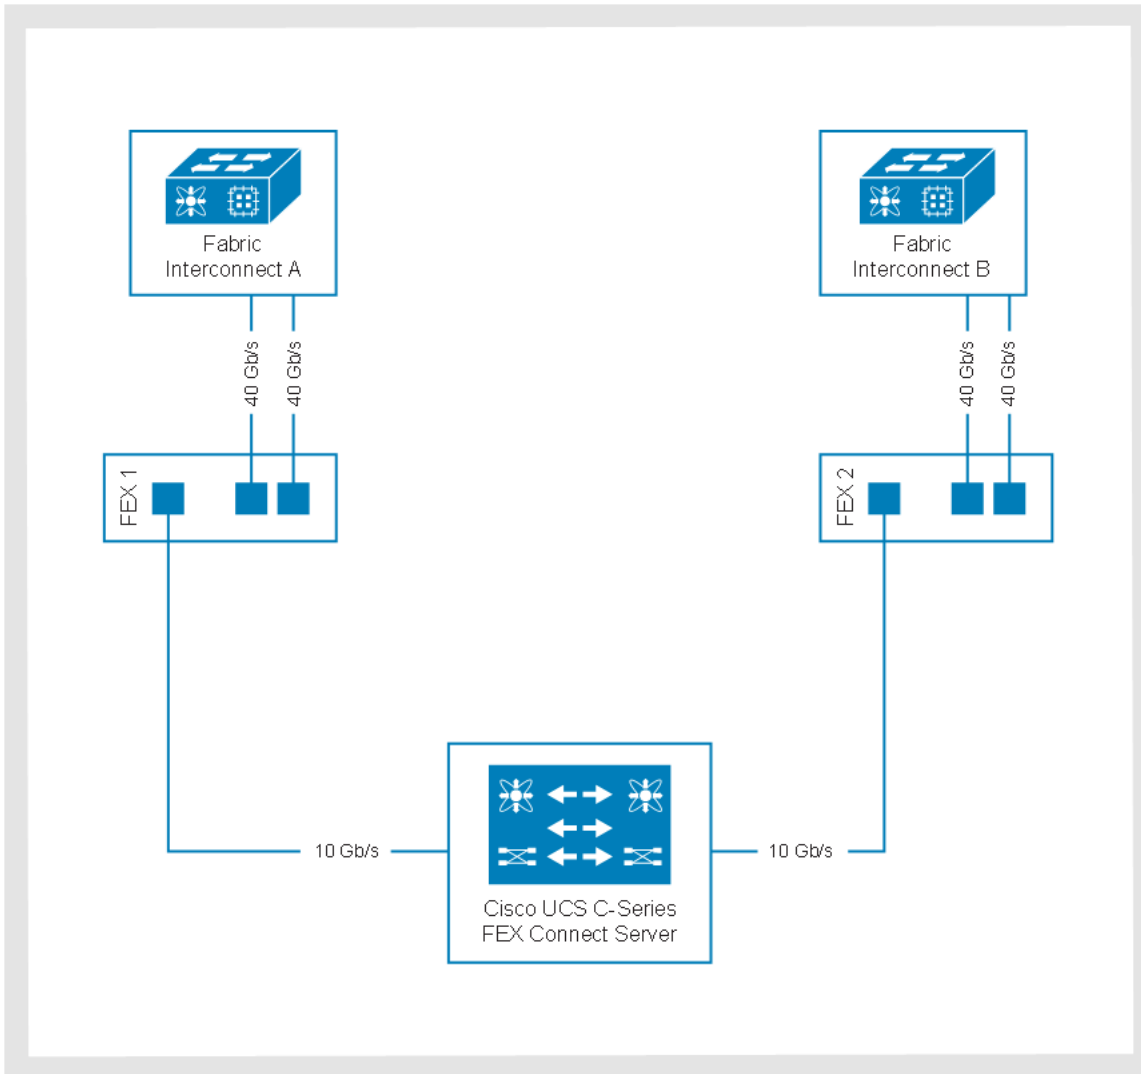
The Cisco Nexus 2348UPQ Fabric Extender allows for a greater number of 10 GbE rack mount servers.

The server connects to the FEX instead of connecting directly into the FI of the system. This method allows for 48 host-capable connections per FEX and six uplink ports.

A Cisco UCS FEX-connected server configuration has the following features:

- Cisco UCS C-Series servers
- Cisco Nexus 2348UPQ 10 GbE FEX
- Two, four or six 40 Gb/s uplinks to the Cisco FI
- Cisco UCS port licensing
- 10 Gb/s per connection per fabric from Cisco UCS M4 servers using the Cisco UCS VIC 1225/1227, while Cisco UCS M5 servers could use VICs 1455/1457 and 1385/1387 or 1495/1497 with supported QSA modules
 - For the Cisco UCS 1455 and 1457 VICs, ports 1 and 2 connect to FEX A. Ports 3 and 4 connect to FEX B.
 - **Cisco UCS 1455/1457 single-link connectivity:** For a two-port connection, port 1 or 2 connects to FEX A and port 3 or 4 connects to FEX B.
 - **Cisco UCS 1455/1457 dual-link connectivity:** For a four-port connection, ports 1 and 2 connect to FEX A and ports 3 and 4 connect to FEX B.

The following figure shows a Cisco UCS Gen 3 configuration with Cisco UCS C-Series FEX connections:



Cisco UCS C-Series Gen 4 Fabric Extender connection

The Cisco Nexus 2232PP Fabric Extender enables for a greater number of 10 GbE rack mount servers.

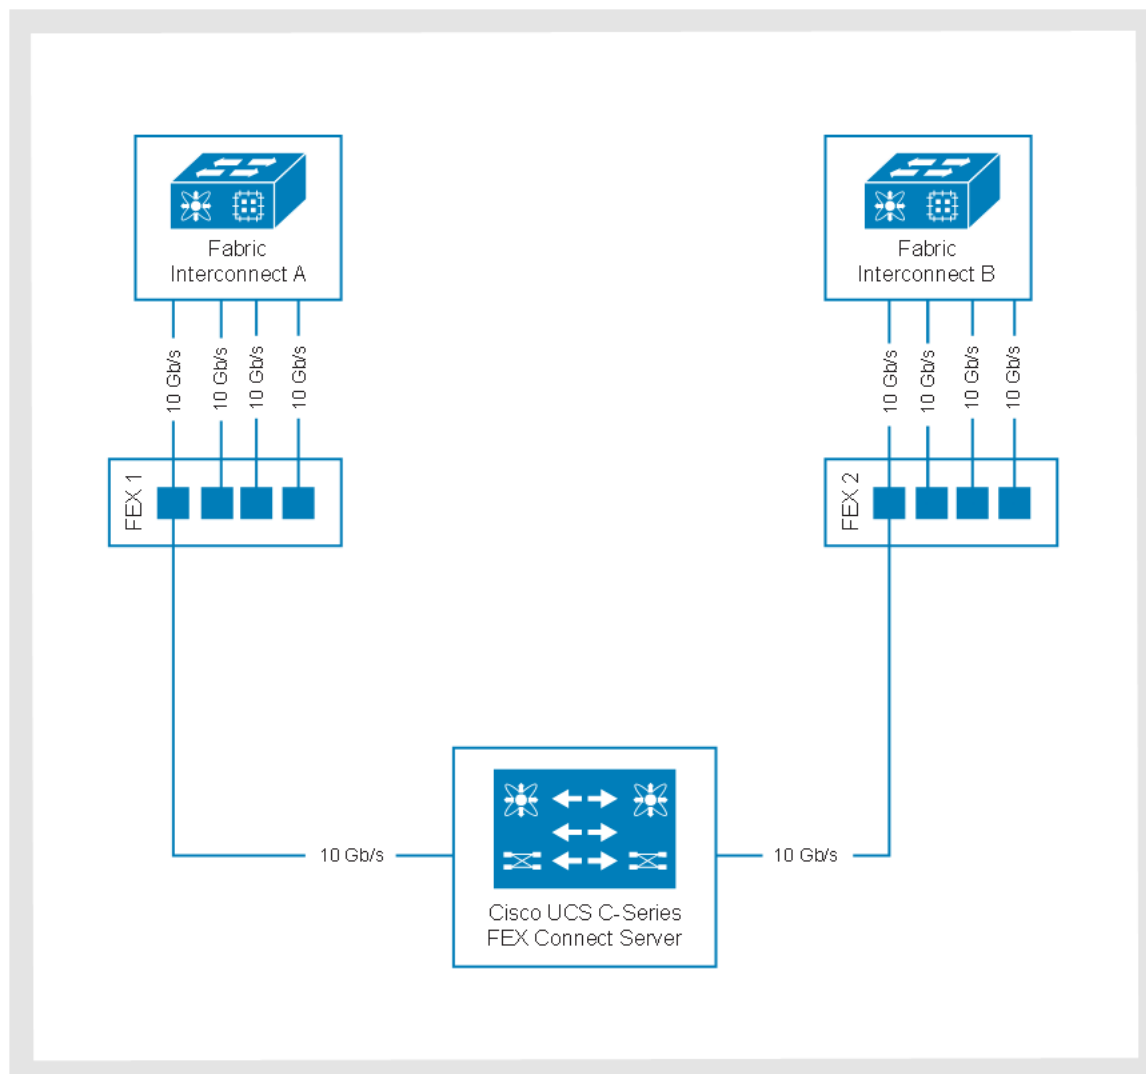
The server connects to the FEX instead of connecting directly into the FI of the system. This method enables for 32 host-capable connections per FEX and six uplink ports.

A Cisco UCS FEX connected server configuration has the following features:

- Cisco UCS C-Series Servers
- Cisco Nexus 2232PP 10 GbE FEX
- Four or eight 10 Gb/s uplinks to the Cisco FI
- Cisco UCS port licensing
- 10 Gb/s per connection per fabric from M4 servers using the Cisco UCS VIC 1225/1227. M5 servers could use VICs 1455/1457 to also achieve 10 Gb/s connectivity.
 - For the Cisco UCS 1455 and 1457 VICs, ports 1 and 2 connect to FEX A. Ports 3 and 4 connect to FEX B.
 - Cisco UCS 1455/1457 single-link connectivity: For a two-port connection, port 1 or 2 connects to FEX A and port 3 or 4 connects to FEX B.

- Cisco UCS 1455/1457 dual-link connectivity: For a four-port connection, ports 1 and 2 connect to FEX A and ports 3 and 4 connect to FEX B.

The following figure shows a Cisco UCS Gen 4 configuration with Cisco UCS C-Series FEX connections:

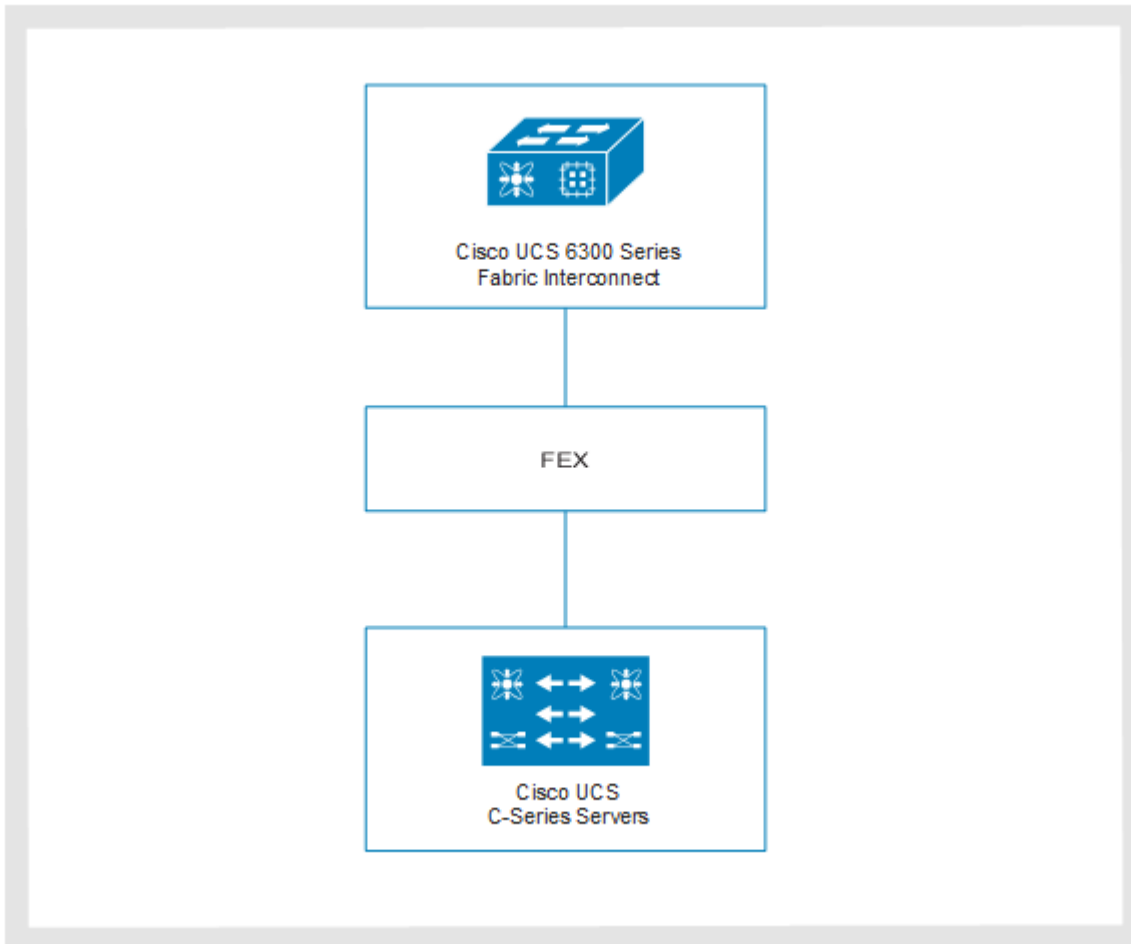


Mixed Cisco UCS B-Series and C-Series servers

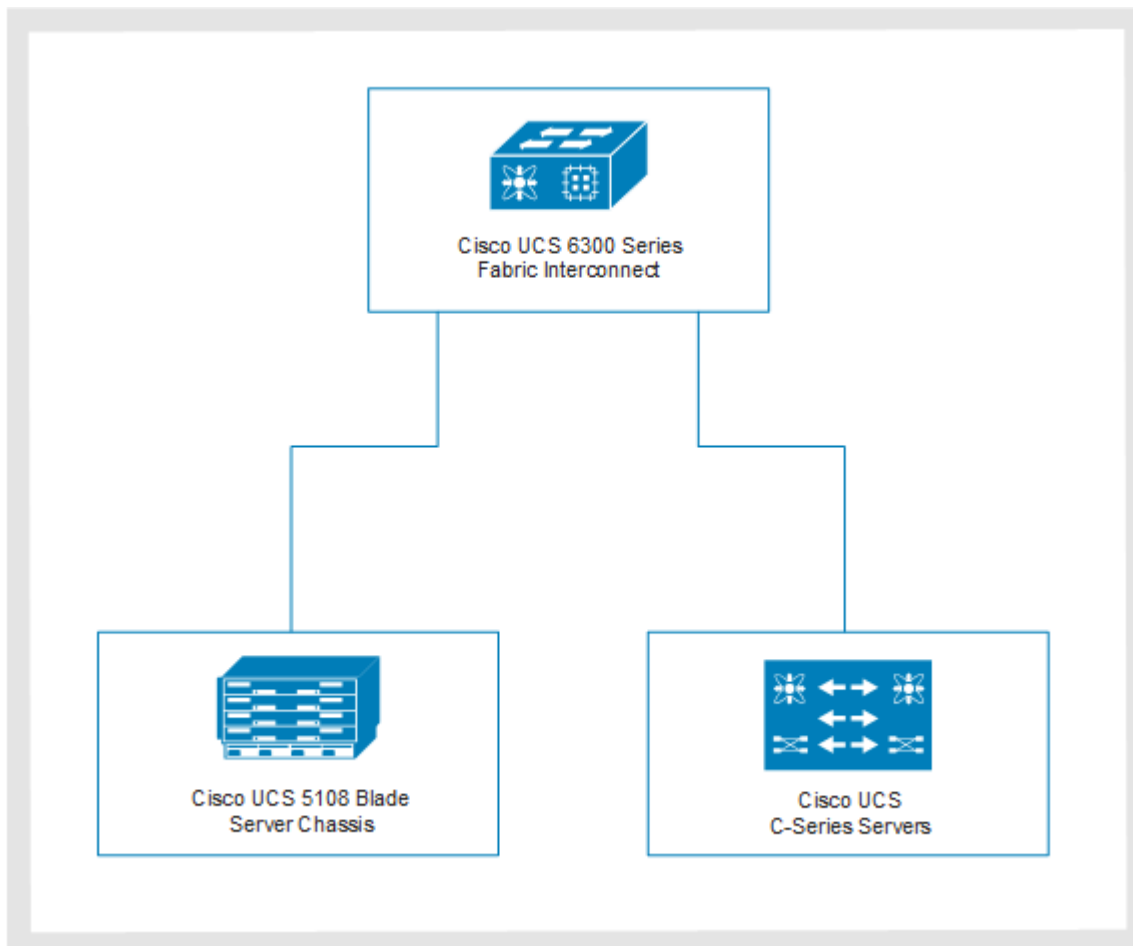
Cisco UCS has a unique architecture that integrates compute, data network access, and storage network access into a common set of components under a single-pane-of-glass management interface.

Cisco UCS fuses access layer networking and servers. The hardware and software components support the Cisco unified fabric, which runs multiple types of data center traffic over a single converged network adapter. By eliminating switching inside a chassis, network access-layer fragmentation is significantly reduced. This simplification reduces the number of switches, cables, adapters, and management points by up to two thirds. All devices in a Cisco UCS domain remain under a single management domain, which remains highly available through the use of redundant components.

The following figure shows a configuration with a FEX connection:



The following figure shows a configuration with a direct connection:



Cisco Trusted Platform Module

The Cisco Trusted Platform Module (TPM) is a computer chip that securely stores artifacts such as measurements, passwords, certificates, or encryption keys that are used to authenticate the Converged System. The Cisco TPM provides authentication and attestation services that enable safer computing in all environments.

The Cisco TPM is available by default in as a component within some Cisco UCS Blade Servers and Rack Servers and is shipped disabled.

Dell EMC supports Cisco TPM hardware but does not support the Cisco TPM functionality. Using Cisco TPM features involves using a software stack from a vendor with significant domain experience in trusted computing. Consult your software stack vendor for configuration and operational considerations relating to the Cisco TPMs.

VMware vSphere 6.7 supports TPM version 2.0.

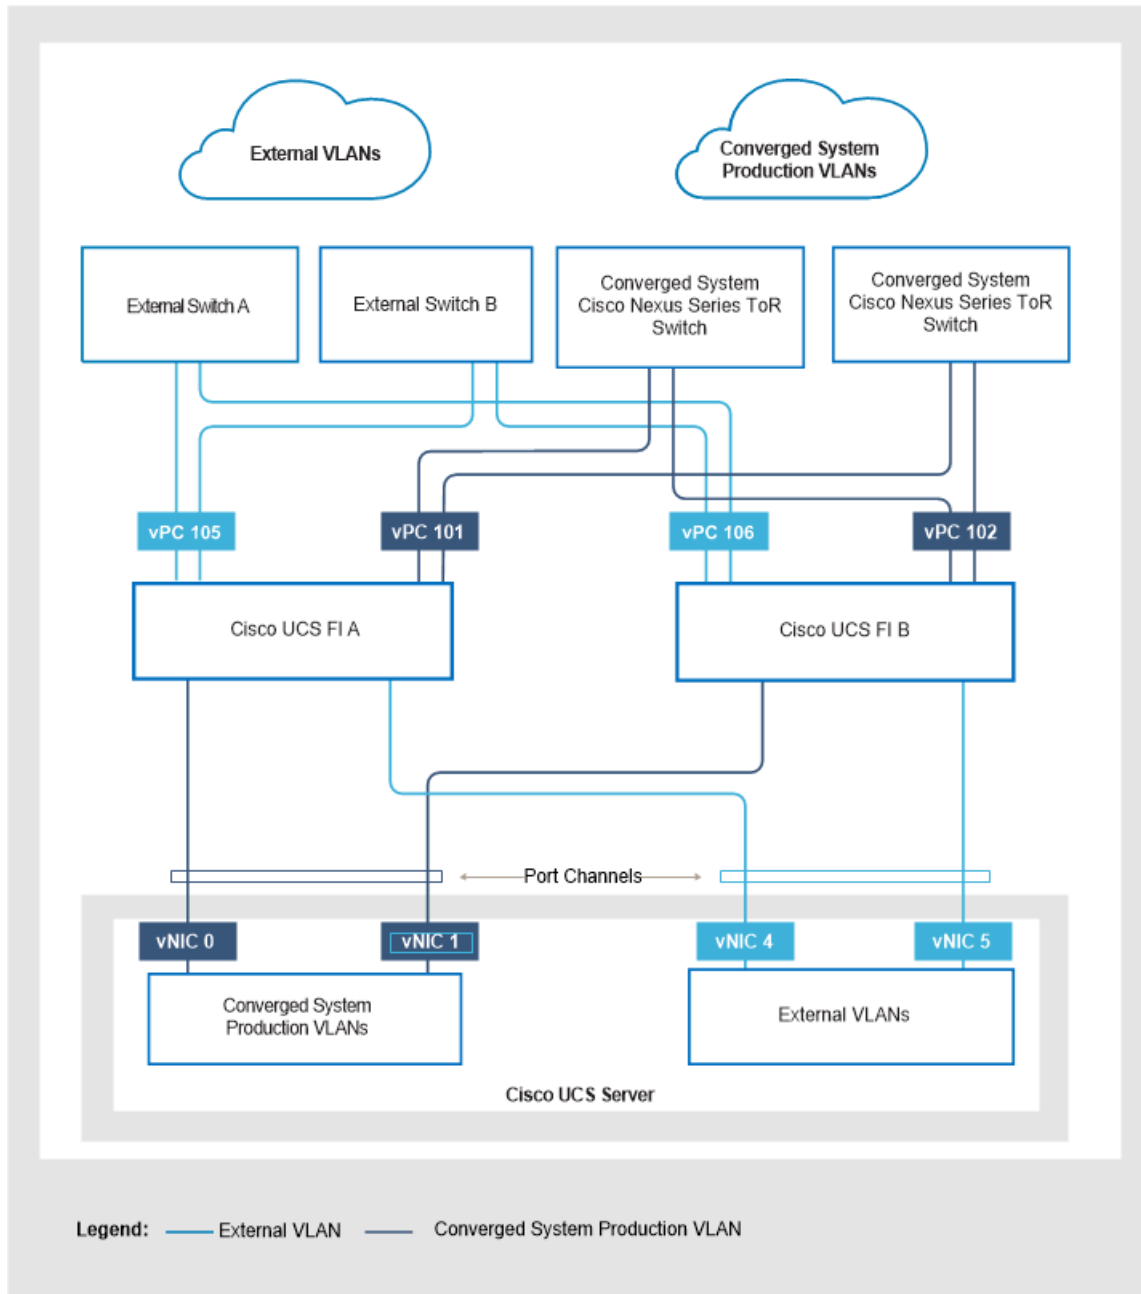
Disjoint Layer 2 configuration

Traffic is split between two or more different networks at the FI in a Disjoint Layer 2 configuration to support two or more discrete Ethernet clouds.

On the Cisco UCS 6300 Series Fabric Interconnects, these can be any of the non-FC 10 Gb/s unified ports or 40 GbE uplink ports not used for standard LAN uplink traffic. On the Cisco UCS 6454 Fabric Interconnects, these can be any of the non FC 10 Gb/s unified ports as well as 40 GbE or 100 GbE Uplink ports not used for standard LAN uplink traffic.

Cisco UCS servers connect to two different clouds. Upstream Disjoint Layer 2 networks allow two or more Ethernet clouds that never connect to be accessed by servers or VMs located in the same Cisco UCS domain.

The following figure provides an example of Disjoint Layer 2 networking into a Cisco UCS domain:



vPCs 101 and 102 are production uplinks that connect to the network layer of the Converged System. vPCs 105 and 106 are external uplinks that connect to other switches.

Disjoint Layer 2 network connectivity can also be configured with an individual uplink on each FI.

Bare metal support policy

Since many applications cannot be virtualized due to technical and commercial reasons, Converged Systems support bare metal deployments, such as non-virtualized operating systems and applications.

While it is possible for Converged Systems to support these workloads (with the following caveats), due to the nature of bare metal deployments, Dell EMC can only provide reasonable effort support for systems that comply with the following requirements:

- Converged Systems contain only Dell EMC published, tested, and validated hardware and software components. The *Release Certification Matrix* provides a list of the certified versions of components for Converged Systems.
- The operating systems used on bare metal deployments for compute components must comply with the published hardware and software compatibility guides from Cisco and Dell EMC.
- For bare metal configurations that include other hypervisor technologies (Hyper-V, KVM, etc.) those hypervisor technologies are not supported by Dell EMC. Dell EMC support is provided only on VMware Hypervisors.

Dell EMC reasonable effort support includes acceptance of customer calls, a determination of whether a Converged System is operating correctly, and assistance in problem resolution to the extent possible.

Dell EMC is unable to reproduce problems or provide support on the operating systems and applications installed on bare metal deployments. In addition, Dell EMC does not provide updates to or test those operating systems or applications. The OEM support vendor should be contacted directly for issues and patches related to those operating systems and applications.

Network layer

The VxBlock System 1000 network layer consists of switches responsible for management, routing, and storage communication.

Data network routing topology

VxBlock System 1000 supports two routing topology configurations, depending on where Layer 3 traffic is routed. One configuration uses Layer 3 routing protocols and the other uses Layer 2 links. The two configurations are explained in the following list:

- The VxBlock 1000 implements the OSPF, EIGRP, or static routing protocol to provide Layer 3 routing services for east-west traffic within the Cisco Nexus 9300 series ToR switches and north-south flows for traffic leaving the VxBlock 1000.
- The VxBlock 1000 is uplinked to the external network with Layer 2 links. In this case, routing for all VxBlock System and external VLANs occurs at the external network. This scenario applies if you need Layer 2 adjacency between endpoints both inside and outside of the VxBlock System. This allows for workload mobility without requiring a network overlay, such as VXLAN, since a subnet can be consumed both inside and outside of the VxBlock System.

This reduces the amount of traffic traversing the uplinks to the external network for VxBlock 1000 inter-VLAN communications and improves overall performance.

Both routing topologies are supported on the management plane as well as the data plane.

The VxBlock System 1000 is configured with a pair of Cisco Nexus 9300 Series TOR switches. If additional connectivity is required, you can add more Cisco Nexus 9300 Series switch pairs.

Management network layer

The management network layer of the VxBlock System 1000 includes at least one pair of Cisco Nexus 31108TC-V management switches.

If more than one pair of Cisco Nexus 31108TV-C switches is required, a pair of Cisco Nexus 3232C Cisco Nexus or 9336C-FX2 switches is used to aggregate management traffic.

The Cisco Nexus management switches provide 1/10 Gb/s connectivity to the management interfaces of the system components and AMP Central, AMP-VX or AMP-3S.

The following table shows management network layer components:

Switch	Description
Cisco Nexus 31108TC-V	<ul style="list-style-type: none"> • 1U appliance • Supports 48 fixed, 100 Mbps/1000 Mbps/10 Gb/s twisted pair connectivity ports and 6 fixed, 40/100 Gb/s QSFP uplink ports for the management layer
Cisco Nexus 3232C	<ul style="list-style-type: none"> • 1U appliance • Supports 32 fixed, 100 Gb/s QSFP28 connectivity ports for aggregating Cisco Nexus 31108TC-V switches • Supports 10/25/40/100 Gb/s SFP+/QSFP28 connections for external uplinks

Switch	Description
Cisco Nexus 9336C-FX2	<ul style="list-style-type: none"> • 1U appliance • Supports 36 fixed, 100 Gb/s QSFP28 connectivity ports for aggregating Cisco Nexus 31108TC-V switches • Supports 10/40/100 Gb/s SFP+/QSFP28 connections for external uplinks

Cisco Nexus 31108TC-V Switch

The base Cisco Nexus 31108TC-V Switches are deployed in pairs to provide 40/100 GbE connections. Each switch provides 48 100 Mbps/1 GbE/10 GbE Base-T fixed ports and six QSFP28 ports to provide 40/100 GbE connections.

The following components connect to the Cisco Nexus 31108TC-V switches using Cat6 connections at 1/10 Gb/s:

- Cisco UCS FI management
- Cisco Nexus 9000 Series ToR switch management
- Cisco MDS SAN switch management
- Storage array management
- IPI Appliance management
- AMP-VX iDRAC
- AMP-VX, AMP-3S, or AMP Central workload
- AMP-VX VMware vSAN
- AMP Central or AMP-VX Data Protection appliance connections
- AMP Central or AMP-3S CIMC
- AMP-3S OOB management iSCSI
- AMP-3S Dell EMC Unity 300 management
- AMP-3S Dell EMC Unity 300 iSCSI
- Data Protection appliance connections

The Cisco Nexus 31108TC-V Switches uplink to the external network using one of the following:

- Cat6 at 1/10 Gb/s
- Multimode fiber at 10/40/100 Gb/s
- Active optical cables (AOC) direct-attach at 100 Gb/s

At least two uplinks per management switch are required.

Cisco Nexus 3232C Switch

The Cisco Nexus 3232C Switch is a management aggregation switch used in the management plane to aggregate pairs of Cisco Nexus 31108TC-V Switches. These switches are required only if the solution requires more than 96 management ports.

The Cisco Nexus 3232C Switch provides:

- 100 Gb/s QSFP ports to aggregate traffic between multiple pairs of Cisco Nexus 31108TC-V Switches
- 100 Gb/s QSFP ports for vPC peer link
- 10/40/100 Gb/s uplinks to the external network using QSFP or QSA adapter

Each Cisco Nexus 31108TC-V Switch has one or two uplinks to each Cisco Nexus 3232C Switch.

Cisco Nexus 9336C-FX2 Switch

The Cisco Nexus 9336C-FX2 Switch is a management aggregation switch used in the management plane to aggregate pairs of Cisco Nexus 31108TC-V Switches. These switches are required only if the solution requires more than 96 management ports.

The Cisco Nexus 9336C-FX2 Switch provides:

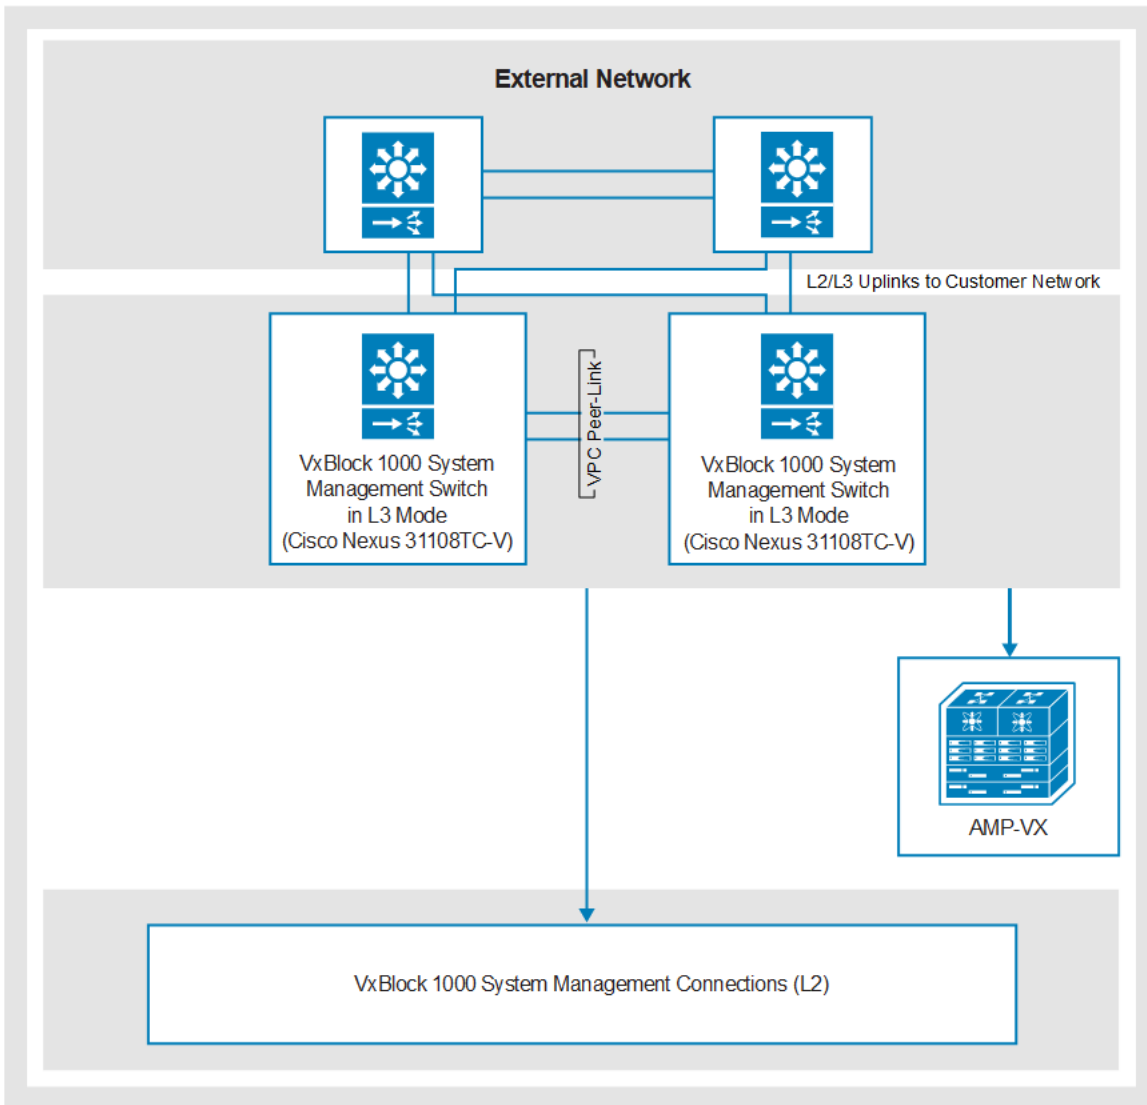
- 100 Gb/s QSFP ports for aggregating traffic between multiple pairs of Cisco Nexus 31108TC-V management switches
- 100 Gb/s QSFP ports for vPC peer link
- 10/40/100 Gb/s uplinks to the external network using QSFP/MPO hydra

Each Cisco Nexus 31108TC-V Switch has one or two uplinks to each Cisco Nexus 9336C-FX2 Switch.

VxBlock System 1000 with no management aggregation layer

The most simple management network deployment for the VxBlock System 1000 consists of a pair of Cisco Nexus 31108TC-V Switches connected directly to the external network.

For a single VxBlock System 1000 with no management aggregation layer, the following figure shows the base management network topology:

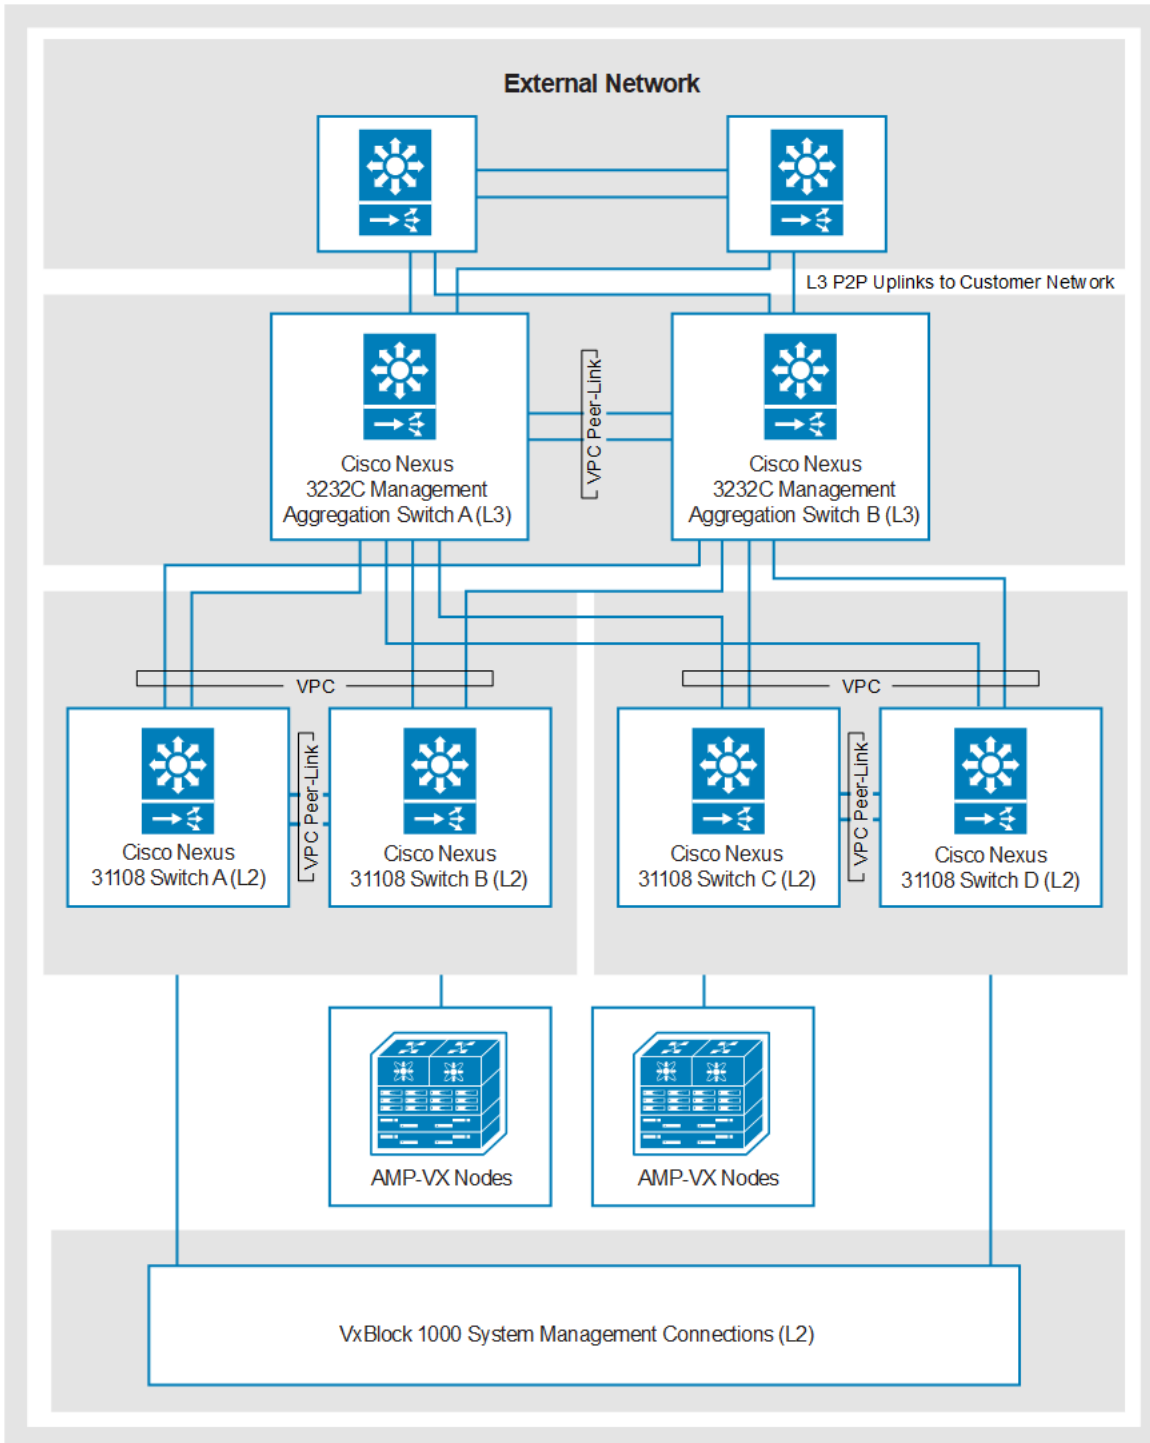


In this scenario, management network SVIs and HSRP addresses can reside on the VxBlock System 1000 management switches (Layer 3) or on the external network (Layer 2).

Single VxBlock System 1000 with a management aggregation layer

This option for a single VxBlock System 1000 introduces the management network to enable a larger scale for the management network.

The following figure shows the management network topology for a single VxBlock System 1000 with the management aggregation layer:



In this scenario, management network Layer 3 SVIs, and HSRP addresses can reside on the VxBlock System 1000 management aggregation switches or on the external network. There is one uplink from each Cisco Nexus 31108TC-V Switch to each Cisco Nexus management aggregation switch at 100 Gb/s.

Multiple VxBlock Systems 1000 with a management aggregation layer

This option for multiple VxBlock Systems 1000 puts the management network on the management aggregation switches.

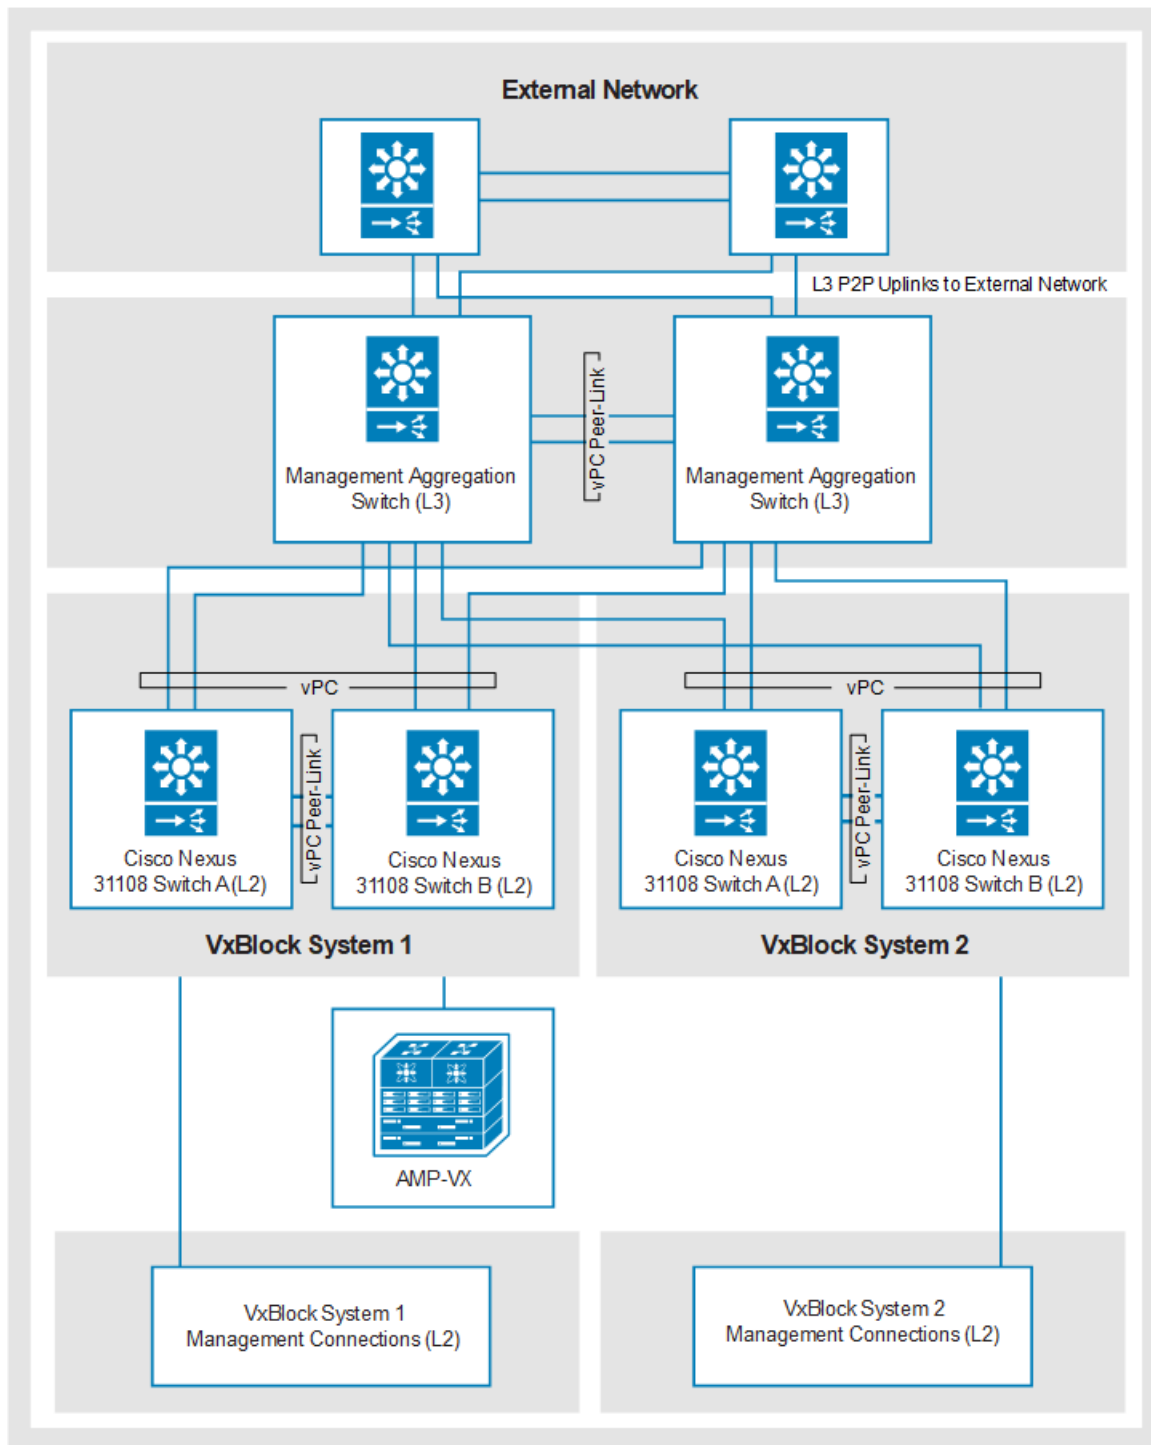
The management aggregation layer contains the following:

- All management network Layer 3 SVIs and HSRP addresses reside on the VxBlock System 1000 management aggregation switches.
- Each connected VxBlock System 1000 has a unique VLAN ID and subnet for OOB management

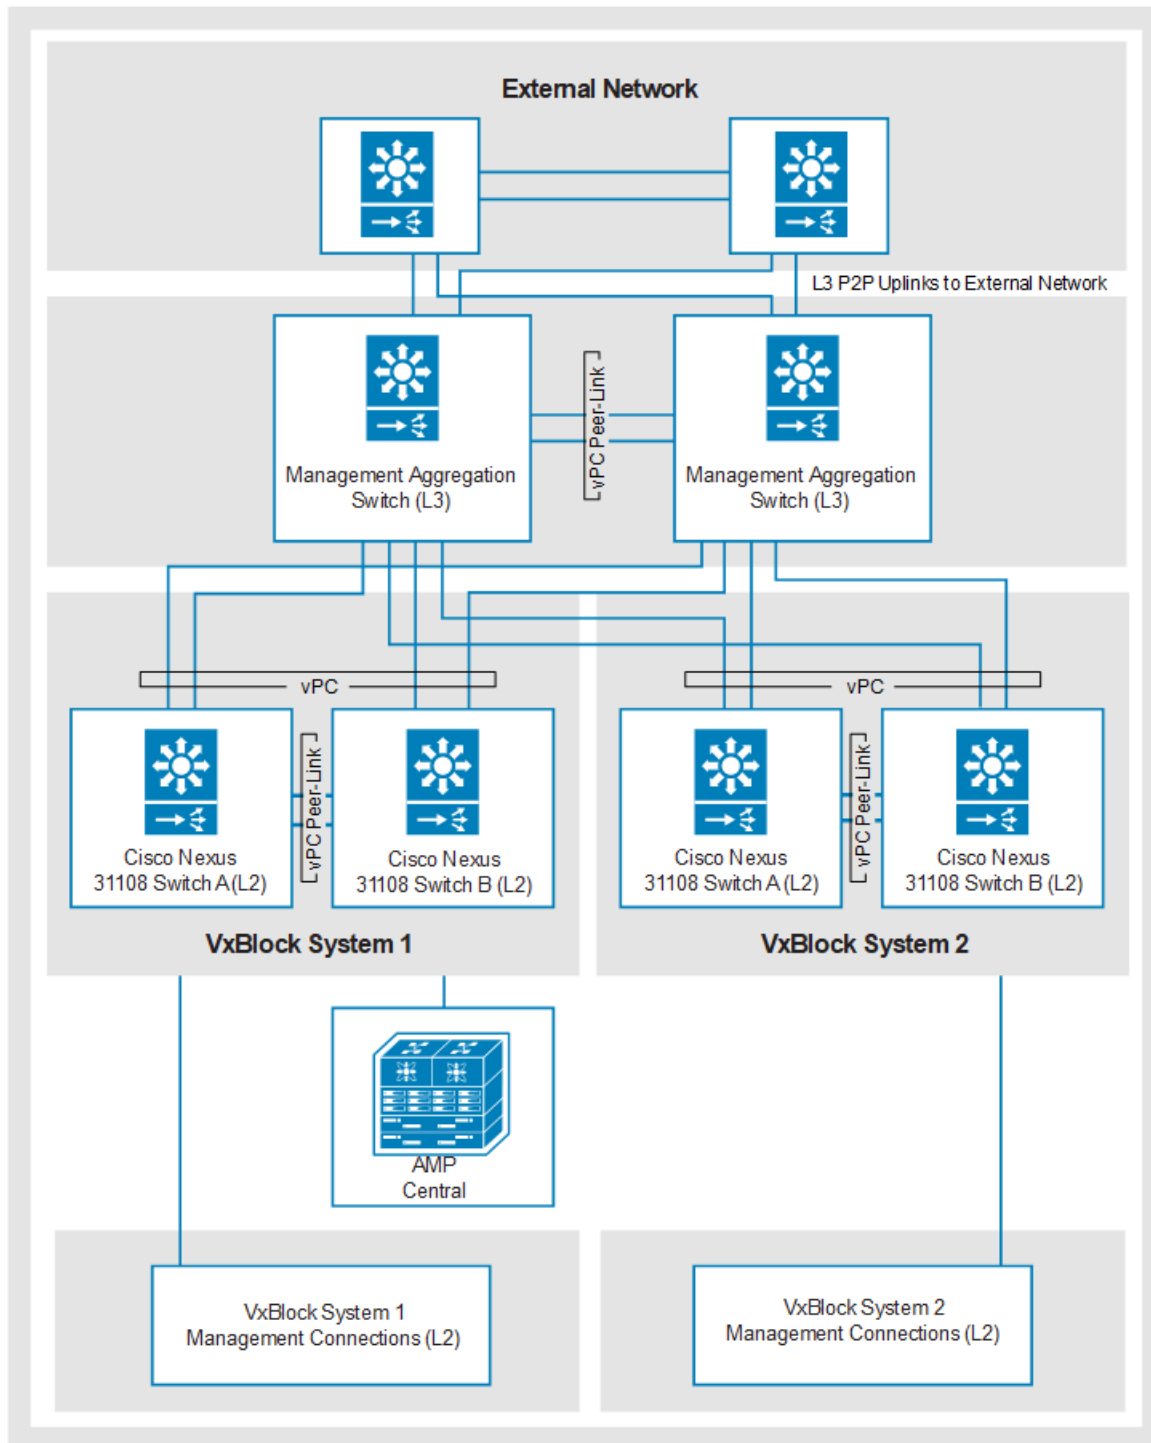
When adding a second VxBlock System to an existing AMP Central or AMP-VX deployment, you may need two Cisco Nexus 9336C-FX2 or Cisco Nexus 3232C management aggregation switches to the second VxBlock System.

During deployment, move all existing SVIs and uplinks for the VxBlock System management switches the Cisco Nexus 9336C-FX2 or Cisco Nexus Cisco Nexus 3232C switch pair.

The following figure depicts AMP-VX with aggregation management:



The following diagram shows AMP Central with management aggregation:



- All management network Layer 3 SVIs and HSRP addresses reside on the VxBlock System 1000 management aggregation switches.
- Each connected VxBlock System 1000 has a unique VLAN ID and subnet for out-of-band management

Software licenses

In deployments where the Cisco Nexus 31108TC-V Switches are the only management switches, if you use Layer 3 uplink configuration, purchase a LAN Enterprise Services license.

Layer 3 protocols require a LAN Enterprise services license. If you choose the Layer 2 uplink model, you do not need to purchase a LAN Enterprise Service license.

All Cisco Nexus 31108TC-V Switches must include a Cisco DCNM LAN license.

For solutions including a management aggregation switch pair, see *Software Licenses* in the Cisco Nexus 9336C-FX2 data plane switch.

Cisco Nexus 93180YC-FX Switch

The Cisco Nexus 93180YC-FX Switch is a data plane that provides 48 10/25 Gb/s SFP+/SFP28 ports and six 40/100 Gb/s QSFP/QSFP28 uplink ports.

The Cisco Nexus 93180YC-FX switch is for solutions that include:

- Up to two Cisco UCS domains with standard uplink bandwidth requirements (160 Gb/s per fabric).
- No Isilon nodes requiring 40 Gb/s connectivity
- No requirement for 40 Gb/s uplinks from the Cisco UCS FIs to the ToR switch

Cisco UCS uplinks

Uplinks from the Cisco UCS 6332-16UP Fabric Interconnects to the Cisco 93180YC-FX Switch use Cisco QSFP to 4x10 Gb/s SFP+ Active Optical breakout cables. Each pair of Cisco UCS FIs consumes 16 SFP+ ports on each Cisco Nexus 93180YC-FX Switch.

Storage array Network-Attached Storage connections

Connectivity from the Network-Attached Storage (NAS) ports on a VxBlock System 1000 storage array to the Cisco Nexus 93180YC-FX Switch is made through a standard 10 Gb/s optical SFP. NAS devices are cross-connected to both switches.

AMP-VX connections

One connection from each AMP-VX node is made to each Cisco Nexus 93180YC-FX Switch with Cisco 10 Gb/s direct-attach cables or 10 Gb/s optics. These connections are used for data-plane traffic, such as connections to the VMware vSphere ESXi Management network from workloads residing on the AMP-VX.

AMP-3S connections

One connection from each AMP-3S server is made to each Cisco Nexus 93180YC-FX Switch with Cisco 10 Gb/s direct-attach cables or 10 Gb/s optics. These connections are used for data-plane traffic, such as connections to the VMware vSphere ESXi Management network from workloads residing on the AMP-3S.

vPC peer link

Two 100 Gb/s ports are used to form the vPC peer-link between the Cisco Nexus 93180YC-FX Switches.

External uplinks

The following table shows the supported connections to the external network:

Bandwidth	Transceiver	Media	Minimum uplinks per switch pair	Maximum uplinks per switch pair	Maximum distance
100 Gb/s	QSFP-100G-SR4-S	OM4 MPO MMF	4	8	100 M

Bandwidth	Transceiver	Media	Minimum uplinks per switch pair	Maximum uplinks per switch pair	Maximum distance
100 Gb/s	QSFP-100G-AOC	Direct-attach Active Optical	4	8	30 M
40 Gb/s	QSFP-40G-SRBD	MMF	4	8	100 M for OM3, 150 M for OM4
40 Gb/s	QSFP-H40G-AOC	Direct-attach Active Optical	4	8	30 M
10 Gb/s	QSFP-40G-SR4	MPO MMF terminated with LC	16	32	100 M for OM3, 150 M for OM4

Integrated Data Protection connections

Most Dell EMC Integrated Data Protection solutions include a pair of dedicated Cisco Nexus 93180YC-FX Switches, which are linked to the Cisco Nexus 93180YC-FX Switches using 10 Gb/s SFP+ optical connections.

For small deployments, direct connectivity from the Integrated Data Protection components to the switches is made with standard 10 Gb/s optical SFP+ based connections to the Cisco Nexus 93180YC-FX Switch. In-band management connections from the Integrated Data Protection components use 1 Gb/s GLC-T SFP+ to Cat6 connections. In deployments with multiple Cisco Nexus 9300 Series Switch pairs, the Integrated Data Protection components can be connected to any Cisco Nexus 9300 Series Switch pair.

Isilon connections

For small Isilon deployments, the ToR Cisco Nexus 93180YC-FX Switches can be used to connect Isilon nodes at 10 Gb/s using SFP-10G-SR transceivers. Isilon deployments that include 40 Gb/s Isilon nodes require Cisco Nexus 93180LC-EX Switches. For large Isilon deployments, extra Cisco Nexus 9300 Series Switch pairs can be added to provide adequate ports. These additional switch pairs are trunked at Layer 2 to the ToR network switches.

Cisco Nexus 93180LC-EX Switch

The Cisco Nexus 93180LC-EX Switch is a data plane switch that provides 24 40 Gb/s QSFP ports and six 40/100 Gb/s QSFP/QSFP28 uplink ports.

This switch can be used for the ToR network and for extra switch pairs to support Isilon connectivity.

The Cisco Nexus 93180LC-EX Switch is for solutions that include:

- Up to five Cisco UCS domains with standard uplink bandwidth requirements (160 Gb/s per fabric)
- Isilon nodes requiring 40 Gb/s connectivity
- Requirements for 40 Gb/s uplinks from the Cisco UCS FIs to the ToR switch

Odd-numbered ports 1, 3, 5, through 23 can be configured as 4x10 Gb/s breakouts to facilitate 10 Gb/s connections to SFP+ based devices. When one of these ports is configured as a breakout, the even-numbered port below it is disabled. Ports 25 and 27 are dedicated as 100 Gb/s vPC peer links, and ports 26 and 28 are disabled. Uplink ports 29 through 32 can be configured as 40 Gb/s, 100 Gb/s, or 4x10 Gb/s breakouts to facilitate connections to external networks.

Cisco UCS uplinks

Uplinks from the Cisco UCS 6332-16UP Fabric Interconnects to the Cisco Nexus 93180YC-FX Switch use:

- Cisco 40 Gb/s QSFP Active Optical cables for the first two Cisco UCS domains
- QSFP optical bi-directional (BiDi) transceivers for Cisco UCS domains three through five

Each pair of Cisco UCS FIs consumes four QSFP ports on each Cisco Nexus 93180LC-EX Switch for a standard uplink topology. Six QSFP ports on each Cisco Nexus 93180LC-EX Switch provide enhanced bandwidth uplink technology.

Storage array NAS connections

A Cisco QSFP-40G-SR4 transceiver provides connectivity from the 10 Gb/s NAS ports on a VxBlock System 1000 storage array to the Cisco Nexus 93180LC-EX Switch.

A combination of MPO-to-LC hydra cables, MPO trunks, and MPO-to-LC cassettes are used to provide fiber connectivity to the endpoint.

The NAS device is populated with 10 Gb/s optical SFP+ modules. A direct-attach QSFP cable that is provided with the node connects to a 40 Gb/s Isilon node. NAS devices are cross-connected to both switches.

AMP-VX/AMP-3S connections

One connection from each AMP node is made to each ToR Cisco Nexus 93180LC-EX Switch using Cisco QSFP-40G-SR4 transceiver.

A combination of MPO-to-LC hydra cables, MPO trunks, and MPO-to-LC cassettes are used to provide fiber connectivity to the endpoint. Fiber connectivity provides four 10 Gb/s LC fiber connections from each QSFP module. These connections are used for data-plane traffic, such as connections to the VMware vSphere ESXi Management network from workloads residing on the AMP.

vPC peer link

Two 100 Gb/s copper direct-attach connections form the vPC peer link between the Cisco Nexus 93180LC-EX Switches.

External uplinks

The following connections to the external network are supported:

Bandwidth	Transceiver	Media	Minimum uplinks per switch pair	Maximum uplinks per switch pair	Maximum distance
100 Gb/s	QSFP-100G-SR4-S	OM4 MPO MMF	4	8	100 M
100 Gb/s	QSFP-40/100G-SRBD	LMMF	4	8	100 M
40 Gb/s	QSFP-40G-SRBD	MMF	4	8	100 M for OM3, 150 M for OM4
40 Gb/s	QSFP-40G-SR4	MPO MMF	4	8	100 M
10 Gb/s	QSFP-40G-SR4	MPO MMF terminated with LC	4 QSFP	8 QSFP	100 M for OM3, 150 M for OM4

Integrated Data Protection connections

Most Dell EMC Integrated Data Protection solutions include a pair of dedicated Cisco Nexus 93180YC-EX Switches, which are uplinked to the Cisco Nexus 93180LC-EX Switches using 40 Gb/s QSFP direct-attach cables.

For small deployments, the Cisco QSFP-40G-SR4 transceiver uses the same connection method as NAS to provide direct connectivity from the data protection components to the switches. This method provides four 10 Gb/s LC fiber connections from each QSFP module. In-band management connections from the Integrated Data Protection components use 1 Gb/s GLC-T SFP+ to Cat6 connections.

Isilon connections

For small Isilon deployments, the ToR Cisco Nexus 93180LC-EX switches can connect Isilon nodes at 40 Gb/s using the QSFP-40G-SR4 transceiver and MPO cabling. 10 Gb/s can be used if implementing the same connection method described for NAS connections. Larger Isilon deployments require extra Cisco Nexus 93180LC-EX Switch pairs.

Cisco Nexus 9336C-FX2 Data Plane Switch (Cisco UCS Gen 3)

The Cisco Nexus 9336C-FX2 Switch provides 36 40/100 Gb/s QSFP28 ports. The switch is used for the ToR network and for more switch pairs to support Isilon connectivity.

Cisco UCS uplinks

Cisco UCS 6332-16UP Fabric Interconnect uplinks to the Cisco Nexus 9336C-FX2 Switch use Cisco 40 Gb/s QSFP Active Optical cables for the first two Cisco UCS domains. Cisco UCS domains three through seven use QSFP optical bi-directional (BiDi) transceivers. Each pair of Cisco UCS FIs consumes the following:

- Four QSFP ports on each Cisco Nexus 9336C-FX2 Switch for a standard uplink topology, or
- Six QSFP ports on each Cisco Nexus 9336C-FX2 Switch for enhanced bandwidth uplink topology

Storage array NAS connections

A Cisco QSFP-40G-SR4 transceiver connects the 10 Gb/s NAS ports on a VxBlock System 1000 storage array to the Cisco Nexus 9336C-FX2 Switch.

A combination of MPO-to-LC hydra cables, MPO trunks, and MPO-to-LC cassettes provide fiber connectivity to the endpoint.

The NAS device is populated with 10 Gb/s optical SFP+ modules. A direct-attach QSFP cable that is provided with the node connects to a 40 Gb/s Isilon node. NAS devices are cross-connected to both switches.

AMP-VX connections

The Cisco QSFP-40G-SR4 transceiver connects each AMP-VX node to each ToR Cisco Nexus 9336C-FX2 Switch.

A combination of MPO-to-LC hydra cables, MPO trunks, and MPO-to-LC cassettes provide fiber connectivity to the endpoint.

This combination provides four 10 Gb/s LC fiber connections from each QSFP module. These connections are used for data plane traffic, such as connections to the VMware vSphere ESXi management network from workloads residing on the AMP-VX.

AMP-3S connections

The Cisco QSFP-40G-SR4 transceiver connects each AMP-3S node to each ToR Cisco Nexus 9336C-FX2 Switch. A combination of MPO-to-LC hydra cables, MPO trunks, and MPO-to-LC cassettes provide fiber connectivity to the endpoint.

This combination provides four 10 Gb/s LC fiber connections from each QSFP module. These connections carry data plane and in-band management traffic for AMP-3S and VxBlock 1000, such as connections to the VMware vSphere ESXi Management network from workloads residing on AMP-3S.

AMP Central connections

The Cisco QSFP-40G-SR4-S Transceiver connects each AMP Central node to each Cisco Nexus 9336C-FX2 ToR Switch. A combination of MPO-to-LC hydra cables, MPO trunks, and MPO-to-LC cassettes provide fiber connectivity to the endpoint. Fiber connectivity provides four 25 Gb/s LC fiber connections from each QSFP module. These connections are used for data plane traffic and in-band management for AMP Central and the VxBlock System. Connections to the VMware vSphere ESXi management network from workloads on AMP Central also use data plane traffic and in-band management.

vPC peer-link

Two 100 Gb/s copper direct-attach connections form the vPC peer link between the Cisco Nexus 93180LC-EX Switches.

External uplinks

The following connections to the external network are supported:

Bandwidth	Transceiver	Media	Minimum uplinks per switch pair	Maximum uplinks per switch pair	Maximum distance
100 Gb/s	QSFP-100G-SR4-S	OM4 MPO MMF	4	16	100 M
100 Gb/s	QSFP-40/100G-SRBD	LMMF	4	16	100 M
40 Gb/s	QSFP-40G-SRBD	MMF	4	16	100 M for OM3, 150 M for OM4
4 x 10 Gb/s	QSFP-40G-SR4	MPO MMF terminated with LC	8 SFP +	32 SFP +	100 M for OM3, 150 M for OM4
40 Gb/s	QSFP-40G-SR4	MPO MMF	4	16	100 M
4 x 10 Gb/s	QSFP-40G-SR4	MPO MMF terminated with LC	8 SFP +	32 SFP +	100 M for OM3, 150 M for OM4

Integrated Data Protection connections

Most Integrated Data Protection solutions include a pair of dedicated Cisco Nexus 93180YC-EX Switches that are uplinked to the Cisco Nexus 9336C-FX2 Switches using 40 Gb/s QSFP direct-attach cables.

For small deployments, direct connectivity from Integrated Data Protection components to the switches is made with Cisco QSFP-40G-SR4 transceiver using the same connection method as the NAS connections. This method provides four 10 Gb/s LC fiber connections from each QSFP module. In-band management connections from the Integrated Data Protection components use 1 Gb/s GLC-T SFP+ to Cat6 connections.

Isilon connections

For small Isilon deployments, the ToR Cisco Nexus 9336C-FX2 Switches can connect Isilon nodes at 40 Gb/s using the QSFP-40G-SR4 transceiver and MPO cabling. 10 Gb/s can be used if implementing the same connection method described for NAS connections. Larger Isilon deployments require more Cisco Nexus 9336C-FX2 Switch pairs.

Software licensing

The Cisco Nexus 9336C-FX2 switch introduces a new tiered licensing model. If you choose the Layer 3 uplink model for the VxBlock 1000, you must purchase a license bundle that includes these services for the ToR switch pair. The license packages that are offered with the VxBlock 1000 are NX-OS Essentials and NX-OS Advanced. Both of these license packages include all L3 protocols, and Cisco DCNM LAN.

If using the Layer 2 uplink, you do not need to purchase a license bundle. You must purchase a Cisco DCNM LAN license for each switch in the solution.

Subordinate switches connected to the ToR Cisco Nexus 9336C-FX2 switch pair do not require a tiered license package, but do require a Cisco DCNM LAN license.

Cisco Nexus 9336C-FX2 Switch (Cisco UCS Gen 4)

The Cisco Nexus 9336C-FX2 Switch is a data plane switch that provides 36 40/100 Gb/s QSFP28 ports. The Cisco Nexus 9336C-FX2 Switch can be used for the ToR network and for extra switch pairs to support Isilon connectivity.

Cisco UCS uplinks

Uplinks from the Cisco UCS 6454 Fabric Interconnects to the Cisco Nexus 9336C-FX2 Switch use:

- Cisco 100 Gb/s QSFP Active Optical cables for the first two Cisco UCS domains
- QSFP optical BiDi-transceivers for Cisco UCS domains three through seven

Each pair of Cisco UCS FIs consumes four QSFP ports on each Cisco Nexus 9336C-FX2 Switch for a standard uplink topology or six QSFP ports on each Cisco Nexus 9336C-FX2 Switch for enhanced bandwidth uplink topology.

Storage array NAS connections

A Cisco QSFP-40G-SR4 transceiver connects the 10 Gb/s NAS ports on a VxBlock System 1000 storage array to the Cisco Nexus 9336C-FX2 Switch.

A combination of MPO-to-LC hydra cables, MPO trunks, and MPO-to-LC cassettes provide fiber connectivity to the endpoint.

The NAS device has 10 Gb/s optical SFP+ modules. A direct-attach QSFP cable that is provided with the node connects to a 40 Gb/s Isilon node. NAS devices are cross-connected to both switches.

AMP-VX connections

The Cisco QSFP-40G-SR4 transceiver connects each AMP-VX node to each ToR Cisco Nexus 9336C-FX2 Switch.

A combination of MPO-to-LC hydra cables, MPO trunks, and MPO-to-LC cassettes provide fiber connectivity to the endpoint.

The configuration provides four 10 Gb/s LC fiber connections from each QSFP module. These connections are used for data-plane traffic, such as connections to the VMware vSphere ESXi Management network from workloads residing on the AMP-VX.

AMP-3S connections

The Cisco QSFP-40G-SR4 transceiver connects each AMP-3S node to each ToR Cisco Nexus 9336C-FX2 Switch. A combination of MPO-to-LC hydra cables, MPO trunks, and MPO-to-LC cassettes provide fiber connectivity to the endpoint. The configuration provides four 10 Gb/s LC fiber connections from each QSFP module. These connections are used for data-plane traffic and in-band management for AMP-3S and VxBlock 1000. For example, the connections link the VMware vSphere ESXi Management network from workloads residing on the AMP-3S.

AMP Central connections

The Cisco QSFP-40G-SR4-S Transceiver connects each AMP Central node to each ToR Cisco Nexus 9336C-FX2 Switch. A combination of MPO-to-LC hydra cables, MPO trunks, and MPO-to-LC cassettes provide fiber connectivity to the endpoint. This provides four 25Gb/s LC fiber connections from each QSFP module. These connections are used for data plane traffic and in-band management for AMP Central and VxBlock System 1000. Connections to the VMware vSphere ESXi Management network from workloads residing on AMP Central also use data plane traffic and in-band management.

vPC peer link

Two 100 Gb/s copper direct-attach connections form the vPC peer link the Cisco Nexus 93180LC-EX Switches.

External uplinks

The following connections to the external network are supported:

Bandwidth	Transceiver	Media	Minimum uplinks per switch pair	Maximum uplinks per switch pair	Maximum distance
100 Gb/s	QSFP-100G-SR4-S	OM4 MPO MMF	4	16	100 M
100 Gb/s	QSFP-40/100G-SRBD	LMMF	4	16	100 M
40 Gb/s	QSFP-40G-SRBD	MMF	4	16	100 M for OM3, 150 M for OM4
40 Gb/s	QSFP-40G-SR4	MPO MMF	4	16	100 M
10 Gb/s	QSFP-10G-SR4	MPO MMF terminated with LC	4 QSFP	8 QSFP	100 M for OM3, 150 M for OM4

Integrated Data Protection connections

Most Integrated Data Protection solutions include a pair of dedicated Cisco Nexus 93180YC-EX Switches, which are uplinked to the Cisco Nexus 9336C-FX2 Switches using 40 Gb/s QSFP direct-attach cables.

For small deployments, direct connectivity from the Integrated Data Protection components to the switches is made with Cisco QSFP-40G-SR4 transceiver using the same connection method as the NAS connections. This provides four 10 Gb/s LC fiber connections from each QSFP module. In-band management connections from the Integrated Data Protection components use 1 Gb/s GLC-T SFP+ to Cat6 connections.

Isilon connections

For small Isilon deployments, the top-of-rack Cisco Nexus 9336C-FX2 Switches can connect:

- Isilon nodes at 40 Gb/s using the QSFP-40G-SR4 transceiver and MPO cabling
- Isilon nodes at 10 Gb/s using the connection method described for NAS connections

Larger Isilon deployments require more Nexus 9336C-FX2 Switch pairs.

Software licensing

The Cisco Nexus 9336C-FX2 switch introduces a new tiered licensing model. If you choose the Layer 3 uplink model for the VxBlock 1000, you must purchase a license bundle that includes these services for the ToR switch pair. The license packages that are offered with the VxBlock 1000 are NX-OS Essentials and NX-OS Advanced. Both of these license packages include all Layer 3 protocols, and Cisco DCNM LAN, which are the features the VxBlock 1000 uses.

If you choose the Layer 2 uplink model, you do not need to purchase a license bundle. However, you must purchase a Cisco DCNM LAN license for each switch in the solution.

Subordinate switches connected to the ToR Cisco Nexus 9336C-FX2 Switch pair do not require a tiered license package, but do require a Cisco DCNM LAN license.

Cisco Data Center Network Manager


Cisco Data Center Network Manager (DCNM) provides a central point of management for Cisco Nexus 9000 Series and Cisco MDS 9000 Series Switches used in the VxBlock System 1000.

Cisco DCNM is deployed as a virtual appliance onto AMP-VX or AMP-3S. A single Cisco DCNM instance manages all connected VxBlock 1000.

Cisco DCNM is divided into two sections:

- LAN for managing Cisco Nexus 9000 Series Switches. This is a mandatory component.
- SAN for managing Cisco MDS Switches. This is also a mandatory component; however, the Cisco DCNM license for the Cisco MDS Switches is optional.

Cisco DCNM provides a graphical user interface for viewing and managing switches, as well as a RESTful API to enable automation.

 **Note:** As of RCM 6.5.10.0 and 6.7.3.0, Cisco DCNM is no longer supported.

Vscale Architecture and ACI readiness

The VxBlock System 1000 is designed with Dell EMC Vscale Architecture and Cisco ACI compatibility in mind.

Each Cisco Nexus 9300 series ToR switch included in a VxBlock 1000 is Vscale Architecture and ACI compatible. Spine-and-leaf data centers can repurpose the existing VxBlock 1000 data network uplinks for spine-to-leaf connectivity with minimal recabling or reconfiguration.

The Dell EMC Vscale Architecture Overview contains additional information.

VMware NSX factory integration

VMware NSX is an option for VxBlock System 1000.

VMware NSX provides a VXLAN network overlay and allows you to implement the full range of SDN features provided by the VMware NSX solution.

Storage network layer

Two Cisco MDS switches/directors make up two separate fabrics to provide 16 or 32 Gb/s of FC connectivity between the compute and storage layer.

The Cisco MDS model selection depends on the total number and speed of SAN ports required to support the compute and storage components.

Connections from the storage components are over 16 Gb/s connections. The number of storage connections varies depending on the storage platform and size.

Cisco UCS Gen 3 FIs provide a FC port channel with one of the following connections:

- Eight 16 Gb/s connections (128 Gb/s bandwidth per fabric)
- 12 connections (192 Gb/s)
- 16 connections (256 Gb/s)

Cisco UCS Gen4 FIs provide a FC port channel with one of the following connections:

- Four 16 Gb/s connections (64 Gb/s bandwidth per fabric)
- Eight 16 Gb/s connections (128 Gb/s)
- Four 32 Gb/s connections (128 Gb/s)
- Eight 32 Gb/s connections (256 Gb/s)

The 32 Gb/s connections require 32 Gb/s SFPs and SAN.

Note: The Cisco MDS 9148S Multilayer Fabric Switch supports eight connections from the Cisco UCS domain and up to two domains.

The Cisco MDS switches provide:

- FC connectivity between compute and storage layer components
- Connectivity for backup and business continuity requirements (if configured)

Note: Inter-Switch Links (ISLs) to the external SAN or between switches are not permitted.

The following table shows SAN network layer components:

Component	Description
Cisco MDS 9148S Multilayer Fabric Switch	<ul style="list-style-type: none"> • 1 U appliance. • Provides 24 to 48 line-rate ports for non-blocking 16 Gb/s throughput. • 12 ports per switch are included in the base license. • One Ports On Demand (POD) license per switch for 12 additional ports is required. • Additional POD licenses can be added to increase the licensed port count to 36 or 48 per switch.

Component	Description
Cisco MDS 9148T Multilayer Fabric Switch	<ul style="list-style-type: none"> • 1 U appliance. • Provides 24 to 48 line-rate ports for non-blocking 32 Gb/s throughput. • 24 ports per switch are included in the base license. • Additional POD licenses can be added to increase the licensed port count to 32, 40, or 48 per switch.
Cisco MDS 9396S Multilayer Fabric Switch	<ul style="list-style-type: none"> • 2 U appliance. • Provides 48 to 96 line-rate ports for non-blocking 16 Gb/s throughput. • 48 ports per switch are included in the base license. • Additional ports can be licensed in increments of 12 ports.
Cisco MDS 9396T Multilayer Fabric Switch	<ul style="list-style-type: none"> • 2 U appliance. • Provides 48 to 96 line-rate ports for non-blocking 32 Gb/s throughput. • 48 ports per switch are included in the base license. • Additional ports can be licensed in increments of 16 ports.
Cisco MDS 9706 Multilayer Director	<ul style="list-style-type: none"> • 9 U appliance. • Provides up to 12 Tb/s front panel FC line rate non-blocking, system level switching. • Dell EMC leverages the advanced 48 port line cards at line rate of 16 Gb/s for all ports. • Dell EMC leverages the advanced 48 port line cards at line rate of 32 Gb/s for all ports. • Consists of two 48 port line cards per director. Two additional 48 port line cards can be added. • Dell EMC requires that four fabric modules are included with all Cisco MDS 9706 Multilayer Directors for an N+1 configuration. • Dell EMC requires that six fabric modules are included with all Cisco MDS 9706 Multilayer Directors with 32 Gb/s Line Cards. • Also includes four PDUs and two supervisors.
Cisco MDS 9710 Multilayer Director	<ul style="list-style-type: none"> • 14 U appliance. • Provides up to 24 Tb/s front panel FC line rate non-blocking, system level switching. • Dell EMC leverages the advanced 48 port line cards at line rate of 16 Gb/s for all ports. • Dell EMC leverages the advanced 48 port line cards at line rate of 32 Gb/s for all ports. • Consists of two 48 port line cards per director. Additional 48 port line cards can be added in pairs (up to 6). • Dell EMC requires that four fabric modules are included with all Cisco MDS 9710 Multilayer Directors for an N+1 configuration. • Dell EMC requires that six fabric modules are included with all Cisco MDS 9710 Multilayer Directors with 32 Gb/s Line Cards. • Dell EMC requires that two additional power supplies are included with all Cisco MDS 9710 Multilayer Directors for a total of eight.

Cisco MDS 9148S Multilayer Fabric Switch

The Cisco MDS 9148S Multilayer Fabric Switch provides 12 to 48 line-rate ports for non-blocking, 16 Gb/s throughput. In the base configuration, 24 ports are licensed (with one POD per switch). Additional ports can be licensed as needed in 12-port increments.

The following table provides core connectivity for the Cisco MDS 9148S Multilayer Fabric Switch per fabric:

Feature	Used ports	Port speeds	Media
Cisco UCS Gen 3 FI uplinks	8 per domain	16 Gb/s	SFP+ and optical cable
Cisco UCS Gen 4 FI uplinks	4 or 8 per domain	16 Gb/s	SFP+ and optical cable
Storage arrays	2 to 40	16 Gb/s	SFP+ and optical cable

 **Note:** Due to TCAM limitations, the Cisco MDS 9148S Switch supports up to 192 servers.

Cisco MDS 9148T Multilayer Fabric Switch

The Cisco MDS 9148T Multilayer Fabric Switch provides 24 to 48 line-rate ports for non-blocking, 32 Gb/s throughput. In the base configuration, 24 ports are licensed. Additional ports can be licensed as needed in 8-port increments.

The following table provides core connectivity for the Cisco MDS 9148T Multilayer Fabric Switch per fabric:

Feature	Used ports	Port speeds	Media
Cisco UCS Gen 3 FI uplinks	8, 12, or 16 per domain	16 Gb/s	SFP+ and optical cable
Cisco UCS Gen 4 FI uplinks	4 or 8 per domain	16 or 32 Gb/s	SFP+ and optical cable
Storage arrays	2 to 40	16 Gb/s	SFP+ and optical cable

Cisco MDS 9396S 16G Multilayer Fabric Switch

The Cisco MDS 9396S 16G Multilayer Fabric Switches provide 48 to 96 line-rate ports for non-blocking, 16 Gb/s throughput.

The base license includes 48 ports. Additional ports can be licensed in 12-port increments. The Cisco MDS 9396S 16 Gb/s Multilayer Fabric Switch is a 96-port fixed switch with no IOM expansion slots.

The following table provides core connectivity for the Cisco MDS 9396S 16 Gb/s Multilayer Fabric Switch per fabric:

Feature	Used ports	Port speeds	Media
Cisco UCS Gen 3 FI uplinks	8, 12, or 16 per domain	16 Gb/s	SFP+ and optical cable
Cisco UCS Gen 4 FI uplinks	4 or 8 per domain	16 Gb/s	SFP+ and optical cable
Storage arrays	2 to 88	16 Gb/s	SFP+ and optical cable

Cisco MDS 9396T 32G Multilayer Fabric Switch

The Cisco MDS 9396T 32G Multilayer Fabric Switches provide 48 to 96 line-rate ports for non-blocking, 16 Gb/s throughput.

The base license includes 48 ports. Additional ports can be licensed in 16-port increments. The Cisco MDS 9396T 32G Multilayer Fabric Switch is a 96-port fixed switch with no IOM expansion slots.

The following table provides core connectivity for the Cisco MDS 9396T 32G Multilayer Fabric Switch per fabric:

Feature	Used ports	Port speeds	Media
Cisco UCS Gen 3 FI uplinks	8, 12, or 16 per domain	16 Gb/s	SFP+ and optical cable
Cisco UCS Gen 4 FI uplinks	4 or 8 per domain	16 or 32 Gb/s	SFP+ and optical cable
Storage arrays	2 to 88	16 Gb/s	SFP+ and optical cable

Cisco MDS 9706 and 9710 Multilayer Directors

Cisco MDS 9706 and 9710 Multilayer Directors are modular directors that provide greater scalability than the fixed SAN switches.

Both models deploy two supervisor modules for redundancy and three fabric modules for intra-director communication. The VxBlock System 1000 requires one additional fabric module for each Cisco MDS 9706 and 9710 Multilayer Director for N+1 redundancy and for full line rate support for all 16 Gb/s FC ports in the directors. All six fabric modules for each Cisco MDS 9706 and 9710 Multilayer Director are required for 32 Gb/s FC line cards.

The Cisco MDS 9706 and 9710 Multilayer Directors are director-class SAN switches with IOM expansion slots for 48-port 16 Gb/s and 32 Gb/s FC line cards. The 9706 has four slots available for line cards and the 9710 has eight. Line cards are added to each director in pairs. A mixture of 16 Gb/s and 32 Gb/s line cards are supported in the same director to support a mixed Cisco UCS environment. Each line card must be ordered in pairs per director. For example, a Cisco MDS 9706 can contain two 16 Gb/s line cards and two 32 Gb/s line cards, but not one of each.

Port licenses are not required for the Cisco MDS 9706 and 9710 Multilayer Directors.

The following table provides core connectivity for the Cisco MDS 9706 and 9710 Multilayer Directors per fabric:

Feature	Used ports	Port speeds	Media
Cisco UCS Gen 3 FI uplinks	8, 12, or 16 per domain	16 Gb/s	SFP+ and optical cable
Cisco UCS Gen 4 FI uplinks	4 or 8 per domain	16 or 32 Gb/s	SFP+ and optical cable
Storage arrays	Cisco MDS 9706: 2–184 Cisco MDS 9710: 2–376	16 Gb/s	SFP+ and optical cable

Storage layer

The VxBlock System 1000 offers several storage arrays, which are detailed in this section.

VMAX All-Flash storage overview

VMAX All-Flash arrays deliver the highest possible flash density by supporting the highest capacity flash drives.

The VMAX All-Flash arrays are flexible to expand performance and capacity to address real-world workloads. With a base capacity of 13 TBu (usable) capacity, the VMAX All-Flash array can provide up to 4.42 PBe (effective) through inline compression. The VMAX All-Flash array can provide up to 6.7 million IOPS of performance.

Array components and specifications

The base component of VMAX All-Flash array is the V-Brick.

VMAX All-Flash components

The V-Brick is comprised of the following components:

- **Director:** Each V-Brick has two redundant directors for storage processing. Each director contains Broadwell-based CPUs, DDR4 memory, and 6 Gb/s or 12 Gb/s SAS back-end connectivity to the SSD drives. There are four slots on each director that are dedicated to front-end connectivity. Those slots can contain FC or Ethernet IO modules. The VMAX 950F and 950FX models have a compression IO module located in one of those four slots. Dell EMC requires two 16 Gb/s FC IO modules for each director. Remaining slots can be used for additional functionality, such as eNAS, SRDF, or ProtectPoint.
- **DAE:** The VMAX 950F/FX DAEs are 6 Gb/s SAS and can contain 120 2.5 inch SSD drives. The VMAX 250F/FX DAEs are 12 Gb/s SAS and can contain 25 2.5 inch SSD drives. Both DAE models support two independent IO channels with automatic failover and fault isolation. There are two DAEs per V-Brick.

Models

The following table shows capacity and performance of the VMAX All-Flash models:

Capacity/performance	VMAX 250F/FX	VMAX 950F/FX
Maximum IOPS	1 million	6.7 million
Initial capacity	11.3 TBu	56.6 TBu
Incremental capacity	11.3 TBu	13.2 TBu
Maximum capacity per array	1.16 PBe	4.42 PBe

The following table shows a component comparison of the VMAX All-Flash models:

Component	VMAX 250F/FX	VMAX 950F/FX
Number of V-Bricks	1 to 2	1 to 8
CPU	Intel Xeon E5-2650-v4 2.2 GHz 12 core	Intel Xeon E5-2697-v4 2.3 GHz 18 core
Number of cores per CPU/Engine/System	12 / 48 / 96	18 / 72 / 576
Cache: Per engine options	512 /1024/2048 GB	1024/2048 GB
Cache: System minimum	512 GB	1024 GB

Component	VMAX 250F/FX	VMAX 950F/FX
Cache: System maximum	4 TB (with 2048 GB per engine)	16 TB (with 2048 GB per engine)
Maximum front-end I/O modules per engine	8	6
DAE	12 Gb/s SAS, 25 drives	6 Gb/s SAS, 120 drives
Maximum drives per V-Brick	50	240
Maximum drives per array	100	1920
Maximum software data movers	4 (3 active plus 1 standby)	8 (7 active plus 1 standby)
Maximum number eNAS I/O modules per software data mover	2	1 (Dell EMC limitation due to slot count)

Storage features

VMAX All-Flash storage offers a number of features, including Integrated Data Protection, encryption, file storage, and replication.

General features

VMAX All-Flash storage arrays have the following features in common:

- Two 16 Gb multimode (MM), FC, four-port IO modules per director (four per engine). The VMAX All-Flash 250F and 250FX have two additional slots per director for front-end connectivity. The VMAX All-Flash 950F and 950FX arrays have one additional slot for front-end connectivity.
- Only 2.5 inch SSD drives are supported for VMAX All-Flash models.
- The number of supported Cisco UCS domains and servers depends on the number of array engines.
- All V-Bricks must contain identical cache types. This standard provides consistent performance throughout the array.
- A single cabinet configuration is available for both models of VMAX All-Flash arrays.
- VMAX All-Flash 950F and 950FX arrays support system bay dispersion of up to 25 meters between system bay 1 and any other system bay.

Integrated Data Protection

The following Integrated Data Protection features are supported for VMAX All-Flash arrays:

- VPLEX
- RecoverPoint and RecoverPoint for VM
- Avamar, Data Domain, ProtectPoint

Encryption

Data At Rest Encryption (D@RE) provides hardware-based, on-array, back-end encryption for VMAX All-Flash arrays by using SAS IO modules that incorporate AES-XTS inline data encryption. These modules encrypt and decrypt data as it is being written to or read from disk. D@RE supports either an internal embedded key manager or an external, enterprise-grade key manager accessible through Key Management Interoperability Protocol (KMIP). The following external key managers are supported:

- SafeNet KeySecure by Gemalto
- IBM Security Key Lifecycle Manager

File storage

Embedded Network Attached Storage (eNAS) is fully integrated into the VMAX All-Flash array. eNAS provides flexible and secure multi-protocol file sharing (NFS 2.0, 3.0, 4.0/4.1), CIFS/SMB 3.0, and multiple file server identities (CIFS

and NFS servers). eNAS uses the Hypermax OS hypervisor to create virtual instances of NAS data movers and control stations on VMAX All-Flash controllers. Control stations and data movers are distributed within the VMAX All-Flash array based on the number of engines and their associated mirrored pair.

By default, VMAX All-Flash arrays are configured with:

- Two control station VMs
- Data mover VMs. The number of data movers varies by array size:
 - The VMAX 250F array can have two (default) or four (maximum configuration that requires two V-Bricks)
 - VMAX 950F array can have two (default), four, six, or eight
- All configurations include one standby data mover.

By default, a single two-port, 10 GbE optical IO module is configured per data mover. A second IO module can be configured for the VMAX All-Flash 250F array. All eNAS IO modules must reside in the same slot on each director on which a data mover resides.

Replication

Native local replication on the VMAX All-Flash array is provided by TimeFinder. This software delivers point-in-time copies of volumes that can be used for backup, decision support, data warehouse refreshes, or any other process that requires parallel access to production data. TimeFinder SnapVX provides snapshot and clone technologies. The following features are available with TimeFinder SnapVX:

- Targetless snapshots – snapshot an entire Storage Group with one command
- Secure snaps – prevents the intentional or unintentional deletion of snapshots
- Cascading snapshots – provides flexibility to present multiple versions of a snapshot

Native remote replication on the VMAX All-Flash array is provided by Symmetrix Remote Data Facility (SRDF). This family of products offers a range of array based disaster recovery, parallel processing, and data migration solutions for VMAX All-Flash arrays. SRDF replicates data between two to four arrays located in the same room, on the same campus, or thousands of kilometers apart. SRDF can utilize Ethernet or FC connections for replication.

SRDF operates in the following modes:

- Synchronous mode (SRDF/S) maintains a real-time copy at arrays located within 200 kilometers. Writes from the production host are acknowledged from the local array when they are written to cache at the remote array.
- Asynchronous mode (SRDF/A) maintains a dependent-write consistent copy at arrays located at unlimited distances. Writes from the production host are acknowledged immediately by the local array, thus replication has no impact on host performance. Data at the remote array is typically only seconds behind the primary site.
- SRDF/Metro makes target devices Read/Write accessible to a host (or multiple hosts in clusters). Hosts write to both the R1 and R2 sides of SRDF device pairs, and SRDF/Metro ensures that each copy remains current and consistent.
- Adaptive copy mode (SRDF/DM) moves large amounts of data quickly with minimal host impact.

SRDF can be tightly integrated with TimeFinder to provide robust solutions to a variety of data replication and disaster recovery needs.

SRDF can be paired with VMware Site Recovery Manager to provide automated storage-based disaster restart operation for VMware environments.

Software licensing

Each VMAX All-Flash model can be purchased with two primary software packages (F and FX). Each package includes software licenses to meet the variety of needs. Extra software licenses can be purchased separately.

The following table identifies licenses that are included in each software package:

Software/feature	F-suite	FX-suite
AppSync Starter Pack	Yes	Yes
Compression	Yes	Yes
CloudArray Enabler	Optional	Yes
D@RE	Optional	Yes
Embedded Management	Yes	Yes
eNAS	Optional	Yes
Non-Disruptive Migration	Yes	Yes
PowerPath	No	75 Licenses
ProtectPoint	Optional	Optional
RecoverPoint	Optional	Optional
SnapVX	Yes	Yes
SRDF	Optional	Yes
SRDF/Metro	Optional	Yes
Storage Analytics	Optional	Optional
Unisphere 360	Optional	Yes
ViPR Suite	Optional	Yes

VMAX storage scale

The number of VxBlock System 1000 hosts supported by the VMAX AF array depends on the number of V-Bricks in the array and the number of 16 Gb/s FC ports available in each. All host IO ports (16 Gb/s FC) are placed in port groups of four ports. Ports are spread across directors, SAN fabrics, and IO modules for redundancy. Each port group can support a maximum of 64 hosts per Dell EMC standard.

The following table indicates the number of hosts supported by a VMAX AF array:

Array model	V-Bricks	Minimum port groups	Hosts supported	Maximum port groups	Hosts supported
250F/FX	1	4	256	8	512
	2	8	512	16	1024
950F/FX	1	4	256	6	384
	2	8	512	12	768
	3	12	768	18	*1152
	4	16	*1024	24	*1536
	5	20	*1280	30	*1920
	6	24	*1536	36	*2304
	7	28	*1792	42	*2688
8	32	*2048	48	*3072	

*Maximum number of hosts supported in a VxBlock 1000 is 800.

Note: A system with the maximum port group count sacrifices functionality of optional features, such as eNAS, SRDF, and ProtectPoint. A VMAX 950F or 950FX with four (or more) V-Bricks supports the maximum number of VxBlock 1000 hosts and still provides optional features.

PowerMax storage overview

PowerMax storage arrays introduce a new family of Dell EMC Non-Volatile Memory Express (NVMe) storage. This new family of storage advances the leadership of Dell EMC All-Flash arrays. PowerMax storage arrays deliver:

- The next generation of low latency measured in microseconds, not milliseconds
- Unparalleled performance and scalability with mission-critical multi-controller storage arrays

With a base capacity of 13 TBu (usable), the PowerMax array can provide up to 4 PBe (effective) through inline compression and deduplication. The PowerMax array can provide up to 6.7 million IOPS of performance.

Array components and specifications

The base component of PowerMax array is the PowerMax brick.

Components

Each PowerMax brick has two redundant directors for storage processing. Each director contains Broadwell-based CPUs, DDR4 memory, and Gen 3 PCIe backend connectivity to the NVMe SSD (dual port) drives.

- Four slots on each director provide front-end connectivity. Those slots can contain FC or Ethernet I/O modules. One compression/dedupe I/O module is in slot 7.
- Two NVMe flash SLICs exist on slots 0 and 6 in PowerMax 2000 models.
- Four NVMe flash SLICs exist on slots 0, 1, 6, and 7 in single-engine PowerMax 8000 models. A compression and dedupe I/O module is in slot 9, which leaves three slots for front-end connectivity, including PowerMax bricks added to the array.
- Three NVMe flash SLICs exist on slots 0, 1, and 6 in multiple-engine PowerMax 8000 models. Dell EMC requires two 16 Gb/s FC I/O modules for each director. Remaining slots may be used for other functionality, such as eNAS, SRDF, or ProtectPoint.
- DAEs can contain 24 2.5-inch PCIe NVMe connected SSD or NAND drives. Each DAE has 16 x 4-lane min-SAS HD PCIe connectors. There are two DAEs per PowerMax brick. PowerMax 8000 models share a second DAE (three DAEs for two PowerMax bricks).

Models

The following tables show capacity and performance of the PowerMax models:

Capacity/Performance	PowerMax 2000	PowerMax 8000
Maximum IOPS	1 million	6.7 million
Maximum capacity per array	1 PBe ¹	4 PBe ¹
Initial capacity per brick (NAND)	13.2 TBu ²	54 TBu
Incremental capacity per brick (NAND)	13.2 TBu ²	13.2 TBu
Initial capacity per brick (SCM)	21.0 TBu ³	21.0 TBu ³
Incremental capacity per brick (SCM)	5.25 TBu ⁴	5.25 TBu ⁴

1. Maximum capacity for each array is based on an overprovisioning ratio of 1.0 for a NAND or mixed SCM/NAND array. All SCM arrays have lower maximum capacities.
2. 13.2TBu base capacity and flash capacity increments are possible with RAID5(3+1) on PowerMax 2000.

3. Base capacities are associated with the selection of an SCM only system. You can add only SCM capacity. SCM is available only with RAID 5(7+1), 1 TB cache engines, and the PRO package.
4. SCM incremental capacity packs are based on RAID 5(7+1). You can have 2.25 Tbu incremental expansion with RAID 5(3+1) on PowerMax 2000.

The following table compares components of the PowerMax models:

Component	PowerMax 2000	PowerMax 8000
Number of PowerMax bricks	1–2	1–8
CPU	Intel Xeon E5-2650-v4 2.5 GHz 12 core	Intel Xeon E5-2697-v4 2.8 GHz 18 core
Number of cores per CPU/engine/system	12/48/96	18/72/5-76
Cache: per engine options	512/1024/2048 GB	1024/2048 GB
Cache: system minimum	512 GB	1024 GB
Cache: system maximum	4 TB (with 2048 GB per engine)	16 TB (with 2048 GB per engine)
Maximum front-end I/O modules per engine	8	8 ¹
Disk array enclosure	Gen 3 PCIe, NVMe SSD 24 drives	Gen 3 PCIe, NVMe SSD 24 drives
Maximum drives per PowerMax brick	48 ²	36 ²
Maximum drives per array	96	288
Maximum software X-Blade	4 (3 active +plus 1 standby) ³	4 (3 active plus 1 standby) ⁴
Maximum number eNAS I/O Modules per software X-Brick	2	2
NVMe drives supported (2.5")	1.92 TB, 3.84 TB, 7.68 TB	1.92 TB, 3.84 TB, 7.68 TB
SCM drives supported (2.5") ⁵	750 GB, 1.5 TB	750 GB, 1.5 TB

1. A single-engine PowerMax 8000 has six I/O modules for the life of the array. Slot 9 is not used.
2. On PowerMax 8000:
 - Each PowerMax brick starts with two 24 slot 2.5-inch DAEs.
 - Each even-numbered PowerMax brick adds a third DAE.
 - The middle DAE is shared between PowerMax brick pairs.
 - The first PowerMax brick has 38 drives including spares, and the second has 34 drives.
 - PowerMax 8000 has 32 usable drives with RAID 5 (7+1) (Default) or RAID 6 (6+2) plus spares.
 - PowerMax 2000 has 40 usable drives with RAID 5 (7+1) (Default) or RAID 6 (6+2) plus spares (2 DAEs). PowerMax 2000 has 44 usable drives with RAID 5 (3+1) plus spares (2 DAEs).
3. Four X-Blade require a minimum of two PowerMax bricks.
4. Six or eight X-Blade are available by RPQ.
5. SCM drives are available with the PowerMaxOS 5978.444.

Storage features

PowerMax storage offers several features, including Integrated Data Protection, encryption, file storage, and replication.

General features

PowerMax storage arrays have the following general features in common:

- Two 16 Gb/s or 32 Gb/s multimode (MM), FC, four-port IO modules per director (four per engine). Two additional slots per director are available for IO modules. The PowerMax 8000 has only one additional slot per director for front-end connectivity in single-engine configurations.
- Only 2.5-inch NVMe SSD drives are supported for PowerMax models.
- The number of supported Cisco UCS domains and servers depends on the number of array PowerMax bricks.
- Intelligent inline compression/deduplication is available.
- A single cabinet configuration is available for both models of PowerMax. For PowerMax 8000, up to four PowerMax bricks are supported in a single cabinet.

Data Protection

The following Integrated Data Protection features are supported for PowerMax:

- VPLEX
- RecoverPoint and RecoverPoint for VM
- Avamar/Data Domain/ProtectPoint

Encryption

Data At Rest Encryption (D@RE) provides hardware-based, on-array, back-end encryption for PowerMax arrays by using SAS IO modules that incorporate AES-XTS inline data encryption. These modules encrypt and decrypt data as it is being written to or read from disk. D@RE supports either an internal embedded key manager, or an external, enterprise-grade key manager accessible through Key Management Interoperability Protocol (KMIP). The following external key managers are supported:

- SafeNet KeySecure by Gemalto
- IBM Security Key Lifecycle Manager

File storage

Embedded Network Attached Storage (eNAS) is fully integrated into the PowerMax array. eNAS provides flexible and secure multiprotocol file sharing (NFS 2.0, 3.0, 4.0/4.1), CIFS/SMB 3.0), and multiple file server identities (CIFS and NFS servers). eNAS uses the PowerMax operating system hypervisor to create virtual instances of NAS Data Movers and Control Stations on PowerMax controllers. Control Stations and Data Movers are distributed within the PowerMax array, depending on the number of engines and their associated mirrored pair.

By default, PowerMax arrays are configured with:

- Two Control Station virtual machines
- Data Mover virtual machines. The number of Data Movers varies by array size:
 - PowerMAX 2000 can have two (default) or four (maximum configuration that requires two PowerMax bricks).
 - PowerMAX 8000 can have two (default) or four (six or eight are available by RPQ.)
- All configurations include one standby data mover.

By default, a single two-port, 10 GbE optical IO module is configured per data mover. A second IO module can be configured for the PowerMax array. All eNAS IO modules must reside in the same slot on each director on which a data mover resides.

Replication

TimeFinder provides native local replication on PowerMax. This software delivers point-in-time copies of volumes for backup, decision support, data warehouse refreshes, or other processes that require parallel access to production data. TimeFinder SnapVX provides snapshot and clone technologies. Following are some of the features available with TimeFinder SnapVX:

- Targetless snapshots enables you to snapshot an entire storage group with one command.
- Secure snaps prevent the intentional or unintentional deletion of snapshots.
- Cascading snapshots present multiple versions of a snapshot.

Symmetrix Remote Data Facility (SRDF) provides native remote replication. This family of products offers disaster recovery, parallel processing, and data migration solutions for PowerMax. SRDF replicates data between two to four arrays in the same room, on the same campus, or thousands of kilometers apart. SRDF can use Ethernet or Fibre Channel connections for replication.

SRDF operates in the following modes:

- Synchronous mode (SRDF/S) maintains a real-time copy at arrays that are located within 200 kilometers. The local array acknowledges writes from the production host when they are written to cache at the remote array.
- Asynchronous mode (SRDF/A) maintains a dependent-write copy at arrays that are at unlimited distances. The local array acknowledges writes from the production host immediately. Replication has no impact on host performance. Data at the remote array is typically only seconds behind the primary site.
- SRDF/Metro makes target devices read/write accessible to a host (or multiple hosts in clusters). Hosts write to both the R1 and R2 sides of SRDF device pairs, and SRDF/Metro ensures that each copy remains current and consistent.
- Adaptive copy mode (SRDF/DM) moves large amounts of data quickly with minimal host impact.

SRDF can be tightly integrated with TimeFinder to provide data replication and disaster recovery.

SRDF can be paired with VMware vCenter Site Recovery Manager to provide automated storage-based disaster restart operation for VMware environments.

Software licensing

Each PowerMax model can be purchased with two primary software packages (Essentials and Pro). Each package includes software licenses to meet the variety of needs. Additional software licenses can be purchased separately.

The following table identifies licenses included in each software package:

Software/Feature	Essentials software package	Pro software package
AppSync Starter Pack	Yes	Yes
Compression/deduplication	Yes	Yes
Data At Rest Encryption	Optional	Yes (not required)
Embedded Management	Yes	Yes
Embedded NAS	Optional	Yes (not required)
PowerPath	No	75 Licenses
ProtectPoint	Optional	Optional
RecoverPoint	Optional	Optional
SnapVX	Yes	Yes
SRDF	Optional	Yes

Software/Feature	Essentials software package	Pro software package
SRDF/Metro	Optional	Yes
Storage Analytics	Yes	Yes
Unisphere 360	Optional	Yes
ViPR Suite	Optional	Yes

PowerMax storage scale

The number of VxBlock System 1000 hosts supported by the PowerMax array depends on the number of PowerMax bricks in the array and the number of 16 Gb FC ports available in each. All host IO ports (16 Gb FC) are placed in port groups of four ports. Ports are spread across directors, SAN fabrics, and IO modules for redundancy. Each port group can support a maximum of 64 hosts.

The following table indicates the number of hosts support by a PowerMax array:

Array model	PowerMax bricks	Minimum port groups	Hosts supported	Maximum port groups	Hosts supported
2000	1	4	256	8	512
	2	8	512	16	1024
8000	1	4	256	6	384
	2	8	512	16	1024
	3	12	768	24	*1536
	4	16	*1024	32	*2048
	5	20	*1280	40	*2560
	6	24	*1536	48	*3072
	7	28	*1792	56	*3584
	8	32	*2048	64	*4096

*Maximum number of hosts supported in a VxBlock 1000 is 800.

Note: To reach the maximum port group count, additional functionality, such as eNAS, SRDF, and ProtectPoint is inhibited. Four (or more) PowerMax bricks with PowerMax 8000 support the maximum number of VxBlock 1000 hosts and still provide optional features.

Dell EMC Unity storage overview

Dell EMC Unity storage arrays are midrange storage solutions that include transactional workloads, such as databases, virtualized servers, and multiple workload consolidations. Dell EMC Unity is a unified solution that supports block and file environments with the flexibility to configure a system, capacity points, and Integrated Data Protection options. Dell EMC Unity storage arrays support integrated block and file storage and are optimized for VMware virtualized applications.

Regardless of the storage protocol implemented, VxBlock Systems 1000 can include cabinet space, cabling, and power to support the hardware. Dell EMC Unity storage includes controllers that eliminate the need for separate hardware for file sharing. The controllers also facilitate moving from block storage to unified storage, avoiding hardware changes.

Dell EMC Unity Hybrid and Unity All-Flash storage arrays provide a scalable hardware design and advanced software capabilities using spinning and flash drives. Dell EMC Unity All-Flash arrays support high-density solid-state disks (SSDs) including 3D NAND triple level cell (TLC) drives.

The following Dell EMC Unity All-Flash storage array models are available:

- 350F
- 450F
- 550F
- 650F

The following Dell EMC Unity Hybrid storage array models are available:

- 300
- 400
- 500
- 600

The Dell EMC Unity All-Flash 350F to 650F storage arrays have the latest system hardware and contain only flash drives. The following table lists the processors and memory per storage processor:

Storage array	Processor (per SP)	Memory (per SP)
350F	Intel E5-2603v4 6-core, 1.7 GHz	48 GB (3 x 16 GB) @ 1866 MHz
450F	Intel E5-2630v4 10-core, 2.2 GHz	64 GB (4 x 16 GB) @ 2133 MHz
550F	Intel E5-2660v4 14-core, 2.0 GHz	128 GB (4 x 32 GB) @ 2400 MHz
650F	Intel E5-2680v4 14-core, 2.4 GHz	256 GB (4 x 64 GB) @ 2400 MHz

The Dell EMC Unity Hybrid storage arrays have the latest system hardware and contain flash and spinning drives. The following table lists the processors and memory per storage processor:

Storage array	Processor (per SP)	Memory (per SP)
300	Intel E5-2603v3 6-core, 1.6 GHz	24 GB (3 x 8 GB) @ 1600 MHz
400	Intel E5-2630v3 8-core, 2.4 GHz	48 GB (3 x 16 GB) @ 1600 MHz
500	Intel E5-2660v3 10-core, 2.6 GHz	64 GB (4 x 16 GB) @ 1866 MHz
600	Intel E5-2680v3 12-core, 2.5 GHz	128 GB (4 x 32 GB) @ 1866 MHz

Array components and specifications

Dell EMC Unity storage arrays contain the disk processor enclosure (DPE) that house the redundant storage processors (SPs).

The DPE provides slots for two SPs and integrated drive slots for 25 2.5-inch drives or 12 3.5-inch drives.

Dell EMC Unity All-Flash models support only the drive carrier in the DPE with 25 Slots for 2.5-inch drives.

Dell EMC Unity Hybrid models support two types of drive carriers in the DPE with either:

- 12 slots for 3.5-inch drives
- 25 slots for 2.5-inch drives


Each SP contains the following components:

- Single socket CPU

Storage layer

- Three to four DDR5 DIMM slots, depending on the model
- Internal battery backup unit (BBU)
- Two integrated 10 GbE BaseT ports (RJ45)
- Two integrated converged network adapter (CNA) ports (SFP+ or Twinax)
 - 10 GbE optical
 - 16 Gb/s FC
- Two integrated 12 Gb/s SAS ports for DAE connections
- Two slots for additional I/O modules (SLICs)
 - Four-port 12 Gb/s SAS (bus expansion)
 - Four-port 16 Gb/s FC
 - Four-port Ethernet 10 GbE optical

The integrated Ethernet ports and the Ethernet I/O modules supply Network Attached Storage (NAS) to associated hosts for file system access. The SPs run virtual NAS servers for file sharing.

 **Note:** VxBlock System 1000 uses the 10 GbE optical option for NAS usage.

DAEs contain individual disk drives and are available in the following configurations:

- Two U model that can house 25 2.5-inch disks
- Three U model that can house 80 2.5-inch disks

Models

The following tables show capacity and performance of the Dell EMC Unity All-Flash storage arrays:

Component	350F	450F	550F	650F
Processor per SP	6 core, 1.7 GHz	10 core, 2.2 GHz	14 core, 2.0 GHz	14 core, 2.4 GHz
Memory (GB per storage array)	96	128	256	512
Maximum drive count	150	250	500	1000
Bus expansion (x4 lane configuration)	NA	NA	4 port, 12 Gb/s SAS	4 port, 12 Gb/s SAS
Bus expansion (x8 lane configuration)	NA	NA	2 port, 12 Gb/s SAS	2 port, 12 Gb/s SAS
LUNs	1000	1500	2000	6000
Pool LUN size (TB)	256	256	256	256
File systems	1000	1500	2000	4000
File system size (TB)	256	256	256	256
Pools	20	30	40	100
NAS servers	90	128	128	256

Models

The following tables show capacity and performance of the Dell EMC Hybrid storage arrays:

Component	300	400	500	600
Processor per SP	6 core, 1.6 GHz	8 core, 2.4 GHz	10 core, 2.6 GHz	12 core, 2.5 GHz
Memory (GB per storage array)	48	96	128	256
Maximum drive count	150	250	500	1000
Maximum FAST cache	800 GB	1.2 TB	3.2 TB	5.4 TB
Bus expansion (x4 lane configuration)	NA	NA	4 port, 12 Gb/s SAS	4 port, 12 Gb/s SAS
Bus expansion (x8 lane configuration)	NA	NA	2 port, 12 Gb/s SAS	2 port, 12 Gb/s SAS
LUNs	1000	1500	2000	6000
Pool LUN size (TB)	256	256	256	256
File systems	500	750	1000	1500
File system size (TB)	256	256	256	256
Pools	20	30	40	100
NAS servers	90	128	128	256

Dell EMC Unity storage features

Dell EMC Unity storage offers several features, including Integrated Data Protection, encryption, file storage, and replication.

Integrated Data Protection

The following Integrated Data Protection features are supported:

- VPLEX
- RecoverPoint and RecoverPoint for VM
- Avamar/Data Domain

Encryption

Dell EMC Unity Data at Rest Encryption (D@RE) prevents unauthorized access to lost, stolen, or failed drives by encrypting all sensitive user data. D@RE uses hardware-based encryption modules in the SAS controllers and 12 Gb/s SAS IO modules. The modules encrypt data written to the drives and decrypt data retrieved from the drives. Because of the controller-based approach, D@RE supports all Dell EMC Unity drive types.

The D@RE feature offers external key management as a security option. If an entire Dell EMC Unity system, including drives, is lost or stolen, external key management prevents unauthorized access.

During system initialization, the D@RE feature is enabled when the license file is installed. Installing the license on the storage array activates D@RE. If the license file includes D@RE, D@RE is permanently enabled on the system and cannot be disabled. Similarly, if the license file does not include D@RE, D@RE is permanently disabled on the system and cannot be enabled. D@RE is enabled or disabled only during initial installation. D@RE does not support enabling encryption on non-D@RE enabled systems.

File storage

Term	Definition
Internal shares/exports	Data from resources inside the VxBlock System 1000. Either VMware vSphere ESXi datastores or file system mounts for guest operating systems
External shares/exports	Data from resources outside of the VxBlock 1000.

If sufficient Ethernet ports can provide separate LACP connections, the same array can provide internal and external file systems.

- Each LACP pair needs a connection to each LAN switch in the VxBlock 1000.
- Both SPs must have the same LACP configurations. For example, if the integrated CNA ports are configured for 10 GbE, Ethernet port 4 connects to LAN switch A and port 5 connects to LAN switch B.
- Both SPs should be cabled the same. The two ports (Ethernet ports 4 and 5) should be an LACP pair. An internal or external NAS server can use the LACP pair. You create an extra LACP pair to allow both internal and external NAS Servers.

Internal and external shares are created with two different NAS servers:

- SP-A hosts one share.
- SP-B hosts one share.

Each server hosts different file systems. The file systems are balanced across the NAS servers. More NAS servers can be created for different VLANs, but always in pairs for balance across the SPs.

Replication

Dell EMC Unity storage arrays provide native support for both synchronous and asynchronous remote replication and support for RecoverPoint to provide remote replication.

For block storage configurations, the VxBlock System 1000 can include RecoverPoint, which provides continuous data protection and remote replication for on-demand protection and recovery. RecoverPoint capabilities include:

- Policy-based management
- Application integration
- Bandwidth reduction

The Dell EMC Unity All-Flash All-Inclusive Base Software package includes RecoverPoint.

To implement RecoverPoint, add two or more RecoverPoint Appliances (RPAs) in a cluster to the VxBlock System 1000. This cluster can accommodate approximately 80 MBps sustained throughput through each RPA.

To ensure proper sizing and performance of an RPA solution, Dell EMC Technical Consultants collect the following information:

- The type of data to be replicated
- Data change rates
- Data growth rates
- Network speeds

FAST Cache and FAST VP for Dell EMC Unity Hybrid arrays

Use these best practices when using FAST Cache and FAST VP on Dell EMC Unity Hybrid arrays.

Use the following guidelines when using FAST Cache on Dell EMC Unity Hybrid arrays:

- Place available flash drives in disk pools before using FAST cache to take advantage of flash for metadata acceleration, snapshots, and replication.

- Use FAST cache for pools that do not contain extreme performance tier drives.
- Separate pools where FAST Cache is and is not active.

Use the following guidelines when using FAST VP on Dell EMC Unity Hybrid arrays:

- Use multi-tier pools where possible and when performance is required. Use a minimum of 5% Extreme Performance (Flash) capacity.
- Create as few pools as possible to support the different I/O performance requirements. Create a separate pool only when the I/O profile for an application differs from the other applications. For example, an application with static data may require a capacity tier-only pool.
- Configure all LUNs in the pool with tiering policy **Start High then Auto-Tier** unless the data profile for the LUN requires the highest available or lowest tier available.
- Schedule FAST VP relocations before backups or nightly batch processing.
- For applications that are continuously active, consider configuring FAST VP relocations to run constantly.
- Maintain at least 10% free capacity in storage pools, so that FAST VP relocations can occur efficiently.

Software licensing

Dell EMC Unity licenses include base and optional software.

All-inclusive base software includes the following:

Management software

- Unisphere: Element Manager
- Unisphere Central: Consolidated dashboard and alerting
- CloudIQ: Cloud-based storage analytics
- Thin provisioning
- Dynamic pools (Not available on Hybrid arrays)
- Traditional pools (Available with All-Flash arrays only through the CLI)
- FAST VP and FAST Cache (Available only with Hybrid arrays)
- Data reduction (Block and File)
- Proactive assist: Configure remote support, online chat, open a service request, and so on
- Quality of Service
- Dell EMC Storage Analytics Adapter for VMware vRealize
- File and Block tiering/archiving to public/private cloud (Cloud Tiering Appliance)

Unified protocols

- File
- Block

Local protection

- Controller-Based Encryption (optional), with internal or external key management
- Local point-in-time copies (Snapshots and thin clones)
- AppSync Basic
- Common Event Enabler; Antivirus Agent, Event Publishing Agent


Remote protection

- Native Asynchronous Block and File Replication

- Native Synchronous Block Replication
- Snapshot Shipping
- RecoverPoint Basic
- RecoverPoint for VMs

Optional software

- AppSync Advanced
- Integrated Data Protection Suite: Backup, archive, and collaboration software
- RecoverPoint Advanced
- PowerPath Migration Enabler
- PowerPath Multipathing
- VPLEX

 **Note:** For more details on software licensing, contact your sales representative.

Storage scale

The number of VxBlock System 1000 hosts that the Dell EMC Unity arrays support depends on the number of 16 Gb/s FC ports available on each SP.

All host IO ports (16 Gb/s FC) are placed in port groups of four ports. Ports are spread across SPs, SAN fabrics, and IO modules for redundancy. Each port group supports a maximum of 64 hosts.

Dell EMC Unity arrays support:

- 4–20 Gb/s FC ports
- 0–16 10 GbE ports

The following table indicates the number of hosts that each Dell EMC Unity array supports:

Array model	FC ports	Port groups	Supported hosts
350F - 650F 300 - 600	4	1	64
	8	2	128
	12	3	192
	16	4	256
	20	5	320

Dell EMC Unity XT storage overview

Dell EMC Unity XT storage arrays are midrange storage solutions that include transactional workloads, such as databases, virtualized servers, and multiple workload consolidations. Dell EMC Unity XT is a unified solution that supports block and file environments with the flexibility to configure a system, capacity points, and Integrated Data Protection options. Dell EMC Unity XT storage arrays support integrated block and file storage and are optimized for VMware virtualized applications. Dell EMC Unity XT storage arrays feature up to two times more IOPS, more memory, and up to 50% more drives than previous Dell EMC Unity models.

Regardless of the storage protocol implemented, VxBlock Systems 1000 can include cabinet space, cabling, and power to support the hardware. Dell EMC Unity XT storage includes controllers that eliminate the need for separate hardware for file sharing. The controllers also facilitate moving from block storage to unified storage, avoiding hardware changes.

Dell EMC Unity XT Hybrid and All-Flash storage arrays provide a scalable hardware design and advanced software capabilities using spinning and flash drives. Dell EMC Unity XT All-Flash arrays support high-density solid-state disks

(SSDs) including 3D NAND triple level cell (TLC) drives. Dell EMC Unity XT All-Flash and Hybrid storage systems with dual-active controller architecture and enterprise-class features are designed for performance, optimized for efficiency with up to 5:1 data reduction.

The following Dell EMC Unity XT All-Flash storage array models are available:

- 380F
- 480F
- 680F
- 880F

The following Dell EMC Unity XT Hybrid storage array models are available:

- 380
- 480
- 680
- 880

The Dell EMC Unity XT All-Flash and Hybrid storage arrays have the latest system hardware. The following table lists the processors and memory per Array:

Storage array	Processors per array	Memory per array
380F/380	2 x Intel CPUs, 12 cores per Array, 1.7 GHz	128 GB
480F/480	2 x dual-socket Intel CPUs, 32 cores per Array, 1.8 GHz	192 GB
680F/680	2 x dual-socket Intel CPUs, 48 cores per Array, 2.1 GHz	384 GB
880G/880	2 x dual-socket Intel CPUs, 64 cores per Array, 2.1 GHz	768 GB

Array components and specifications

Based on the powerful family of Intel Xeon processors, Dell EMC Unity XT storage systems implement an integrated architecture for block and file, with concurrent support for native NAS, iSCSI, and Fibre Channel protocols. Each system leverages dual storage processors, full 12 GB SAS back end connectivity and the Dell EMC patented multicore architected operating environment to deliver unparalleled performance and efficiency. Additional storage capacity is added via Disk Array Enclosures (DAEs).

Dell EMC Unity XT storage arrays contain the disk processor enclosure (DPE) that houses the redundant storage processors (SPs) and integrated drive slots for 25 2.5-inch drives. The Dell EMC DPE for the Unity XT 380 and 380F arrays is considered to be enclosure 0 on Bus 0 and The DPE for the 480, 480F, 680, 680F, 880, and 880F storage arrays is considered to be enclosure 0 on Bus 99.

- DAEs contain individual disk drives and are available in the following configurations
 - 2-U model that can house 25 2.5-inch disks
 - 3-U model that can house 80 2.5-inch disks
 - 3-U model that can house 15 3.5-inch disks

Each SP contains the following components:

- Dual socket CPUs
- Three to four DDR5 DIMM slots, depending on the model
- Internal battery backup unit (BBU)

For Dell EMC Unity XT 380/380F storage arrays each SP has the following connectivity components:

- Two integrated 1 GbE BaseT ports (RJ45) for management and service

Storage layer

- Two integrated 10 GbE BaseT ports (RJ45) for front-end connectivity
- Two integrated converged network adapter (CNA) ports (SFP+ or Twinax)
 - 10 GbE optical
 - 16 Gb/s FC
- Two integrated 12 Gb/s SAS ports for DAE connections
- Two slots for additional I/O modules (SLICs)
 - Four-port 16 Gb/s FC
 - Four-port Ethernet 10/25 GbE optical

The integrated Ethernet ports and the Ethernet I/O modules supply Network Attached Storage (NAS) to associated hosts for file system access. The SPs run virtual NAS servers for file sharing.

For Dell EMC Unity XT 480, 480F, 680, 680F, 880, and 880F storage arrays, each SP has the following connectivity components:

- Two integrated 1 GbE BaseT ports (RJ45) for Mgmt and Service
- Two integrated 12 Gb/s SAS ports for DAE connections
- Optional 4-Port Mezz card for front-end connectivity
 - Four-port 10/25GbE Optical with 10GbE or 25GbE SFPs (Can mix SFP's on the same card)
 - Four-port 10GbE BaseT
- Two slots for additional I/O modules (SLICs)
 - Four-port 16 Gb/s FC
 - Four-port Ethernet 10/25 GbE optical
 - Four-port 12 Gb/s SAS (bus expansion)

The optional Mezzanine card or Ethernet I/O modules supply Network Attached Storage (NAS) to associated hosts for file system access. The SPs run virtual NAS servers for file sharing.

VxBlock System 1000 uses the 10 GbE or 25 GbE optical option for NAS usage.

The following table shows capacity and performance of the Dell EMC Unity XT All-Flash and Dell EMC Unity XT Hybrid storage arrays:

Component	380 and 380F	480 and 480F	680 and 680F	880 and 880F
Processors per array	12 cores, 1.7 GHz	32 cores, 1.8 GHz	48 cores, 2.1 GHz	64 cores, 2.1 GHz
Memory (GB per storage array)	128	192	384	768
Maximum drive count	500	750	1,000	1,500
Bus expansion (x4 lane configuration)	N/A	2x 4 Lane 12 Gb/s SAS	2x 4 Lane 12 Gb/s SAS	2x 4 Lane 12 Gb/s SAS
Bus expansion (x8 lane configuration)	N/A	2 x 8 lane 12 Gb/s SAS	2 x 8 lane 12 Gb/s SAS	2 x 8 lane 12 Gb/s SAS
LUNs	1,000	1,500	2,000	6,000
Pool LUN size (TB)	256	256	256	256
File systems	1,000	1,500	2,000	4,000
File system size (TB)	256	256	256	256
Pools	20	30	40	100

Component	380 and 380F	480 and 480F	680 and 680F	880 and 880F
NAS servers	90	128	129	256

Dell EMC Unity XT storage features

Dell EMC Unity XT storage offers several features, including Integrated Data Protection, encryption, file storage, and replication.

Integrated Data Protection

The following Integrated Data Protection features are supported:

- VPLEX
- RecoverPoint and RecoverPoint for VM
- Avamar/Data Domain

Encryption

Dell EMC Unity XT Data at Rest Encryption (D@RE) prevents unauthorized access to lost, stolen, or failed drives by encrypting all sensitive user data. D@RE uses hardware-based encryption modules in the SAS controllers and 12 Gb/s SAS IO modules. The modules encrypt data written to the drives and decrypt data retrieved from the drives. Because of the controller-based approach, D@RE supports all Dell EMC Unity XT drive types.

The D@RE feature offers external key management as a security option. If an entire Dell EMC Unity system, including drives, is lost or stolen, external key management prevents unauthorized access.

During system initialization, the D@RE feature is enabled when the license file is installed. Installing the license on the storage array activates D@RE. If the license file includes D@RE, D@RE is permanently enabled on the system and cannot be disabled. Similarly, if the license file does not include D@RE, D@RE is permanently disabled on the system and cannot be enabled. D@RE is enabled or disabled only during initial installation. D@RE does not support enabling encryption on non-D@RE enabled systems.

File storage

Term	Definition
Internal shares/exports	Data from resources inside the VxBlock System 1000. Either VMware vSphere ESXi datastores or file system mounts for guest operating systems
External shares/exports	Data from resources outside of the VxBlock 1000.

If there are enough Ethernet ports to provide separate LACP connections, the same array can provide internal and external file systems.

- Each LACP pair needs a connection to each LAN switch in the VxBlock 1000.
- Both SPs must have the same LACP configurations. For example, if the integrated CNA ports are configured for 10 GbE, Ethernet port 4 connects to LAN switch A and port 5 connects to LAN switch B.
- Both SPs should be cabled the same way. The two ports (Ethernet ports 4 and 5) should be an LACP pair. An internal or external NAS server can use the LACP pair. Create an extra LACP pair to enable both internal and external NAS Servers.

Create internal and external shares with two different NAS servers:

Each server hosts different file systems. The file systems are balanced across the NAS servers. More NAS servers can be created for different VLANs, but always in pairs for balance across the SPs.

Replication

Dell EMC Unity XT storage arrays provide native support for both synchronous and asynchronous remote replication and support for RecoverPoint to provide remote replication.

For block storage configurations, the VxBlock System 1000 can include RecoverPoint, which provides continuous data protection and remote replication for on-demand protection and recovery.

RecoverPoint capabilities include:

- Policy-based management
- Application integration
- Bandwidth reduction

The Dell EMC Unity All-Flash All-Inclusive Base Software package includes RecoverPoint.

To implement RecoverPoint, add two or more RecoverPoint Appliances (RPAs) in a cluster to the VxBlock System 1000. This cluster can accommodate approximately 80 MB/s sustained throughput through each RPA.

To ensure proper sizing and performance of an RPA solution, Dell EMC Technical Consultants collect the following information:

- The type of data to be replicated
- Data change rates
- Data growth rates
- Network speeds

FAST Cache and FAST VP for Dell EMC Unity XT Hybrid arrays

Use these best practices when using FAST Cache and FAST VP on Dell EMC Unity XT Hybrid arrays.

Use the following guidelines when using FAST Cache on Dell EMC Unity XT Hybrid arrays:

- Place available flash drives in disk pools before using FAST cache to take advantage of flash for metadata acceleration, snapshots, and replication.
- Use FAST cache for pools that do not contain extreme performance tier drives.
- Separate pools where FAST Cache is and is not active.

Use the following guidelines when using FAST VP on Dell EMC Unity XT Hybrid arrays:

- Use multi-tier pools where possible and when performance is required. Use a minimum of 5% Extreme Performance (Flash) capacity.
- Create as few pools as possible to support the different I/O performance requirements. Create a separate pool only when the I/O profile for an application differs from the other applications. For example, an application with static data may require a capacity tier-only pool.
- Configure all LUNs in the pool with tiering policy **Start High then Auto-Tier** unless the data profile for the LUN requires the highest available or lowest tier available.
- Schedule FAST VP relocations before backups or nightly batch processing.
- For applications that are continuously active, consider configuring FAST VP relocations to run constantly.
- Maintain at least 10% free capacity in storage pools, so that FAST VP relocations can occur efficiently.

Software licensing

Dell EMC Unity XT licenses include base and optional software.

All-inclusive base software includes the following:

Management software

- Unisphere: Element Manager
- Unisphere Central: Consolidated dashboard and alerting

- CloudIQ: Cloud-based storage analytics
- Thin provisioning
- Dynamic pools (Not available on Hybrid arrays)
- Traditional pools (Available with All-Flash arrays only through the CLI)
- FAST VP and FAST Cache (Available only with Hybrid arrays)
- Data reduction (Block and File)
- Proactive assist: Configure remote support, online chat, open a service request, and so on
- Quality of Service
- Dell EMC Storage Analytics Adapter for VMware vRealize
- File and Block tiering/archiving to public/private cloud (Cloud Tiering Appliance)

Unified protocols

- File
- Block

Local protection

- Controller-Based Encryption (optional), with internal or external key management
- Local point-in-time copies (Snapshots and thin clones)
- AppSync Basic
- Common Event Enabler; Antivirus Agent, Event Publishing Agent Remote protection
- Native Asynchronous Block and File Replication
- Native Synchronous Block Replication
- Snapshot Shipping
- RecoverPoint Basic
- RecoverPoint for VMs Optional software
- AppSync Advanced
- Integrated Data Protection Suite: Backup, archive, and collaboration software
- RecoverPoint Advanced
- PowerPath Migration Enabler
- PowerPath Multipathing
- VPLEX

For more details on software licensing, contact your sales representative.

Storage scale

The number of VxBlock System 1000 hosts that the Dell EMC Unity XT arrays support depends on the number of 16 Gb/s FC ports available on each SP.

All host IO ports (16 Gb/s FC) are placed in port groups of four ports. Ports are spread across SPs, SAN fabrics, and IO modules for redundancy. Each port group supports a maximum of 64 hosts.

Dell EMC Unity XT 380 and 380F arrays support:

- 4–20 16 Gb/s FC ports
- 0–16 10/25 GbE ports

Dell EMC Unity XT 480, 680, 880, 480F, 680F, and 880F arrays support:

Storage layer

- 8–16 16 Gb/s FC ports
- 0–16 10 GbE BaseT ports
- 0–16 10/25 GbE ports

Dell EMC Unity XT 480, 480F, 680, 680F, 880, and 880F require a four-port 12 Gb/s SAS backend I/O module to reach the maximum drive count.

The following table indicates the number of hosts that each Dell EMC Unity XT array supports:

Array model	FC ports	Port groups	Number of supported hosts
380 and 380F	4	1	64
	8	2	128
	12	3	192
	16	4	256
	20	5	320
Unity XT 480, 680, 880, 480F, 680F, and 880F	8	2	128
	16	4	256

Isilon storage overview

Isilon storage is a scale-out NAS storage solution that delivers increased performance for file-based data applications and workflows from a single file-system architecture.

Isilon storage provides scale-out capacity for use as NFS and SMB CIFS shares within the VMware vSphere VMs.

Depending on the specific configuration, Isilon storage may contain the following key hardware components:

Node type	Processor (per node)	Memory (per node)
F800	Intel E5-2697Av4 16-core, 2.6 GHz	256 GB
H600	Intel E5-2680v4 14-core, 2.2 GHz	256 GB
H500	Intel E5-2630v4 10-core, 2.2 GHz	128 GB
H400	Intel D-1527 4-core, 2.2 GHz	64 GB
A200	Intel D-1508 2-core, 2.2 GHz	16 GB
A2000	Intel D-1508 2-core, 2.2 GHz	16 GB

The following Cisco Nexus switches provide front-end connectivity:

- Cisco Nexus 93180YC-FX
- Cisco Nexus 93180LC-EX
- Cisco Nexus 9336C-FX2

The Isilon back-end Ethernet switches have the following features:

- 10 GbE 24 port
- 10 GbE 48 port
- 10 GbE 96 port (2 x 48 port leaf modules)
- 10 GbE 48 port leaf module

- 40 GbE 32 port
- 40 GbE 64 port (2 x 32 port leaf modules)
- 40 GbE 32 port leaf module

Note: Leaf modules are applicable only to chassis type and in 10 GbE over 48 nodes and 40 GbE over 32 nodes.

Array components and specifications

Isilon All-Flash, hybrid, and archive nodes are contained within a chassis that holds four nodes.

Each chassis contains the following components:

- Four compute slots per chassis
 - Compute suitcase
 - Single-socket CPU
 - Four DDR4 DIMM slots
 - Front-end 10 GbE or 40 GbE optical (depending on node type)
 - Back-end 10 GbE or 40 GbE optical (depending on node type)
 - Single onboard 1 GbE Ethernet
 - DB-9 serial connection
 - Up to 2 x 2.5-inch SSD for Cache
 - Up to 2 x M.2 vault disks
 - Dual 1100 W or 1450 W PSUs
- Drive sleds of 20 2.5-inch drives each
 - F800
 - H600
- Drive sleds of 20 3.5-inch drives each
 - H500
 - H400
 - A200
 - A2000

Note: VxBlock System 1000 does not use 1 GbE connections in each compute suitcase.

The following models are supported:

Component	F800	H600	H500	H400	A200	A2000	
Processor per node	16 core 2.6 GHz	14 core 2.4 GHz	10 core 2.2 GHz	4 core 2.2 GHz	2 core 2.2 GHz	2 core 2.2 GHz	
Memory per node (GB)	256	256	128	64	16	16	
	1.6 TB SSD: 96 TB	600 GB SAS: 72 TB	2 TB Hard drive: 120 TB	2 TB Hard drive: 120 TB	2 TB Hard drive: 120 TB	10 TB Hard drive: 800 TB	

Component	F800	H600	H500	H400	A200	A2000	
Chassis capacity	3.2 TB SSD: 192 TB	1.2 TB SAS: 144 TB	4 TB Hard drive: 240 TB	4 TB Hard drive: 240 TB	4 TB Hard drive: 240 TB		
	15.4 TB SSD: 924 TB		8 TB Hard drive: 480 TB	8 TB Hard drive: 480 TB	8 TB Hard drive: 480 TB		
Front-end networking	2 x 10 GbE or 40 GbE	2 x 10 GbE or 40 GbE	2 x 10 GbE or 40 GbE	2 x 10 GbE	2 x 10 GbE	2 x 10 GbE	
Back-end networking	2 x 40 GbE	2 x 40 GbE	2 x 40 GbE	2 x 10 GbE	2 x 10 GbE	2 x 10 GbE	

Storage features

Isilon supports encryption, file storage, and replication features.

Encryption

The OneFS system is available as a cluster that is composed of Isilon OneFS nodes that contain only self-encrypting drives (SEDs). The system requirements and management of data-at-rest on self-encrypting nodes are identical to the nodes without self-encrypting drives. Clusters of mixed node types are not supported.

Self-encrypting drives store data on an Isilon cluster designed for data-at-rest encryption (D@RE). D@RE on self-encrypted drives occurs when data stored on a device is encrypted to prevent unauthorized data access. All data written to the storage device is encrypted when it is stored, and all data read from the storage device is decrypted when it is read. The stored data is encrypted with a 256-bit data AES encryption key and decrypted in the same manner. OneFS controls data access by combining the drive authentication key with on-disk data-encryption keys. SED options are not included.

File storage

Isilon OneFS combines the three layers of traditional storage architectures (file system, volume manager, and data protection) into one unified software layer, creating a single intelligent distributed file system that runs on an Isilon storage cluster.

VxBlock 1000 configures the two front-end interfaces of each node in an LACP port channel. The front-end interfaces are then used using SmartConnect to load balance share traffic across the nodes in the cluster depending on the configuration.

Replication

Isilon OneFS leverages the SyncIQ licensed feature for replication. SyncIQ is an application that enables you to flexibly manage and automate data replication between two Isilon clusters. SyncIQ delivers unique, highly parallel replication performance that scales with the dataset to provide a solid foundation for disaster recovery. SyncIQ can send and receive data on every node in the Isilon cluster so replication performance is increased as your data grows.

Software licensing

Isilon OneFS comes with various licensable features that expand the functionality of the Isilon cluster.

The following table lists the Isilon features:

Feature	Details
CloudPools	Cloud tiering
Security hardening	Cluster security (STIG, and so on)
HDFS	Hadoop file system protocol
Isilon Swift	OneFS Swift object API
SmartConnect	Cluster connection load balancing
SmartDedupe	Data deduplication
SmartLock	WORM data immutability
SmartPools	Data tiering
SmartQuotas	Quota management
SnapshotIQ	File system snapshots
SyncIQ	Cluster asynchronous replication
Isilon Gen 6 Nodes	Current generation of Isilon cluster hardware

Perpetual and subscription models


The following licensing models are available:

Subscription model	Type	Software
Perpetual	Basic bundle	SmartConnect, SnapshotIQ
	Enterprise Bundle	SmartConnect, SnapshotIQ, SmartQuotas
	Enterprise Advanced Bundle	SmartConnect, SnapshotIQ, SmartQuotas, SyncIQ, SmartPools
Subscription	OneFS Essentials Subscription	SmartConnect, SnapshotIQ, SmartQuotas
	OneFS Advanced Subscription	All software except CloudPools
	OneFS CloudPools Third Party Subscription	CloudPools for third party

Isilon storage scale

The number of Isilon nodes that are supported depends on the number of 10 GbE or 40 GbE ports available in the system.

All node front-end ports (10 GbE or 40 GbE) are placed in LACP port channels. The front-end ports for each of the nodes are connected to a pair of redundant network switches.

 **Note:** For Isilon 8.1, the maximum Isilon configuration requires two pairs of ToR switches.

The following table indicates the number of nodes that are supported for Isilon 8.1:

Model	Node scalability	Capacity scale
F800	4–96 @ 10 GbE 4–48 @ 40 GbE	96 TB per chassis
H600	4–96 @ 10 GbE 4–48 @ 40 GbE	72 TB per chassis

Model	Node scalability	Capacity scale
H500	4–96 @ 10 GbE 4–48 @ 40 GbE	120 TB per chassis
H400	4–96 @ 10 GbE	120 TB per chassis
A200	4–96 @ 10 GbE	120 TB per chassis
A2000	4–96 @ 10 GbE	800 TB per chassis

XtremIO storage overview

The XtremIO Storage Array is an all-flash system based on a scale-out architecture. The system uses building blocks called X-Brick building blocks which are clustered together using an InfiniBand backbone Remote Direct Memory Access (RDMA) network. The system operation is managed using a stand-alone dedicated Linux-based server, called the XtremIO Management Server (XMS).

The XtremIO array architecture is specifically designed to deliver the full performance potential of flash, while linearly scaling all resources such as CPU, RAM, SSDs and host ports.

Due to its content-aware storage architecture, XtremIO provides:

- Even distribution of data blocks, leading to maximum performance and minimal flash wear
- Even distribution of metadata
- No data or metadata hotspots
- Easy setup and no tuning
- Advanced storage functionality, including inline data deduplication and compression, thin provisioning, advanced data protection (XDP), snapshots, and more

Array components and specifications

The base component of the XtremIO storage array is the X-Brick. Each X-Brick consists of two redundant active/active storage controllers. The new multidimensional scalability of X2 enables you to scale up by adding as few as six SSDs at a time to a single X-Brick. You can also scale out with more X-Brick building blocks in any odd or even increments.

The management platform used to manage the XtremIO clusters is the XtremIO Management Server (XMS). The XMS is a VM deployed on the AMP-VX, AMP Central, or AMP-3S. It is not in the data path. An XMS can manage up to 8 XtremIO X2 clusters. X1 clusters can be managed from the XMS provided it is running version 6.0.1-30 or later.

The XMS controls and manages the system, including:

- Monitoring system health and events
- Monitoring system performance
- Maintaining a performance statistics history database
- Providing webUI and CLI services to clients
- Providing RESTful API interfaces

Models

X2 is enhanced with newer Intel CPUs and expanded memory to optimize performance and enable consistent sub-millisecond latency. XtremIO X2-R and X2-S X-Brick building blocks are available. X2-S is for use cases such as virtual desktops with lower capacity requirements and benefit highly from compression and deduplication. X2-R is for workloads that do not benefit from compression and deduplication. These applications require higher physical capacity such as databases.

The following table describes both X-Brick building blocks:

Component	X2-S	X2-R
Number of X-Brick building blocks	1–4	1–4
CPU	Intel® Xeon® processors 2 x 12 core 2.5 GHz Haswell 24 cores per controller (48 per X-Brick)	Intel® Xeon® processors 2 x 12 core 2.5 GHz Haswell 24 cores per controller (48 per X-Brick)
Number of cores per controller/X-Brick/ cluster	24 / 48 / 192	24 / 48 / 192
Minimum number of storage controllers	2	2
Maximum number of storage controllers	8	8
Memory: per X-Brick options	384 GB	1024 GB
Minimum FC ports (16 Gb/s) per X-Brick	4	4
Maximum 16 Gb/s FCs per cluster	16	16
DAE quantity and speed	four 12-Gb/s SAS, 72 drives	four 12-Gb/s SAS, 72 drives
DAE SSD sizes	400 GB/1.92 TB	1.92 TB/3.84 TB
Maximum SSD drives per X-Brick	72	72
Minimum SSD enclosures	1	1
Maximum SSD enclosures	4	4
Maximum Ethernet port for management (10 Gb/s)	2	2
Maximum InfiniBand switches	2	2
InfiniBand switches	12 ports (56 Gb/s)	36 ports (56 Gb/s)

Storage features

Describes the features of XtremIO X2 storage.

XtremIO X2 includes:

- Top load 72-drive DAE for higher density
- Scale-Up and Scale-Out architecture
- HTML GUI
- Inline deduplication and compression
- XtremIO Integrated Copy Data Management for snapshots
- Thin Provisioning
- D@RE

Each X-Brick is composed of:

- One 2U Disk Array Enclosure (DAE) containing:
 - Up to 72 SSDs

- Two redundant power supply units
- Two redundant SAS interconnect models
- 12 Gb/s SAS Connectivity
- Two 1U Storage Controllers
Each Controller includes:
 - Two redundant power supply units (PSUs)
 - Two 16 Gb/s FC ports
 - Two 56 Gb/s InfiniBand HCA
 - One 10 Gb/s management port
 - NVRAM for power loss protection

XtremIO X2 storage scale

Describes the relationship between hosts, FC ports, and X-Bricks.

The number of VxBlock System 1000 hosts supported by the XtremIO X2 arrays depends on the number of X-Bricks in the cluster.

Each X-Brick contains four 16 Gb/s ports and each host should connect to at least four ports in the cluster.

The host ports are spread between storage controllers, X-Bricks, and SAN fabrics for redundancy.

Each group of four ports can support a maximum of 64 hosts in line with Dell EMC standards.

The following table indicates the number of hosts supported by X2 arrays:

X-Brick count	FC ports	Supported hosts
1	4	64
2	8	128
3	12	192
4	16	256

Scale-out architecture

An XtremIO storage system can include a single X-Brick or multiple X-Bricks. Currently both the X2-R and X2-S support up to four X-Bricks

With clusters of two more X-Bricks, XtremIO utilizes two redundant 56 Gb/s InfiniBand HCA ports for an ultra-low latency back-end connectivity between storage controllers. The X2-S includes a 12 port InfiniBand switch while the X2-R includes a 36 port InfiniBand switch for future expansion.

Scale-up architecture

More capacity can be added to an existing configuration without adding compute resources. The minimum number of drives in a DAE is 18. Afterwards disks can be added in packs of 6 until a total of 36 are reached. After 36 have been added, the next disk addition must include 18 additional disks for a total of 54. Afterwards, disks may again be added in packs of six until the DAE is full at 72 disks.

Supported Integrated Data Protection options

XtremIO Virtual Copies (XVC) are created at a particular point in time, allowing users to access the data when needed, even when the source volume has been changed.

VPLEX VS2 and VS6 are supported.

RecoverPoint Gen6 is supported.

Multiple storage array overview

The VxBlock System 1000 offers multiple storage arrays natively without the addition of a Converged Technology Extension.

You can add as many storage arrays or platforms as the connectivity model allows.

Block storage

You can combine multiple storage arrays for block-level storage.

The following best practices are implemented when combining multiple storage arrays and storage array platforms for block level storage:

- Each port group per array is spread across the SAN fabric.
- Hosts do not share array workloads when performance is required.
- VMware vSphere ESXi advanced settings are configured.
- When mixing storage array platforms, common settings are used.
- Host level and/or array level maximum boundaries are maintained per array.

 **Note:** Mixing storage array platforms may affect performance.

File storage

You can combine multiple storage arrays for file-level storage.

The following best practices are implemented when combining multiple storage arrays and storage array platforms for file-level storage:

- Storage array Ethernet ports are spread across the LAN switches.
- Each pair of ports is configured in an LACP port channel.
- VMAX, Dell EMC Unity, and Dell EMC Unity XT provide access to VMware vSphere ESXi datastore.
- Guest access to file shares is supported by VMAX, Dell EMC Unity, Dell EMC Unity XT, and Isilon.
- Host-level and/or array-level maximum boundaries are maintained per array.

Manage and monitor multiple storage arrays

Multiple storage arrays have management and monitoring features.

Unisphere for multiple VMAX arrays

Unisphere 360 enables you to access the data center from any device, including tablets and mobile devices. Unisphere 360 consolidates and simplifies management for VMAX storage arrays. Managing a data center with multiple storage systems can be a large-scale project on its own. Although Unisphere for VMAX allows multiple arrays to be managed from the same instance, the maximum number of devices and arrays are limited. Unisphere 360 supports up to 200 systems covering VMAX All-Flash, VMAX3 and VMAX. A single window view with access to all VMAX storage arrays allows you to manage, monitor, and plan at the array level or for the entire data center.

Unisphere and CloudIQ for multiple Dell EMC Unity and Dell EMC Unity XT arrays

Unisphere Central obtains aggregated status, alerts, host details, performance and capacity metrics, and storage usage information from the systems. Unisphere Central allows administrators to take a quick look at their entire storage environment and rapidly access storage systems that are in need of attention or maintenance.

Storage layer

CloudIQ is a no cost software-as-a-service (SaaS) cloud-based storage analytics application. CloudIQ uses near real-time intelligence, proactive monitoring, and predictive analytics to deliver comprehensive health scores for Dell EMC Unity and Dell EMC Unity XT storage systems from the cloud.

Management

Use VxBlock Central to manage and monitor VxBlock Systems in a data center.

With VxBlock Central, you can:

- View the health and RCM compliance of multiple VxBlock Systems.
- View charts of key performance indicators (KPI) for one or more components or elements.
- Download software and firmware components to maintain compliance with the current RCM.
- Track real-time information regarding critical faults, errors, and issues affecting VxBlock Systems.
- Configure multisystem AD integration and map AD Groups to VxBlock Central roles.
- Set up compute, storage, networks, and PXE services, manage credentials, and upload ISO images for server installation.
- Monitor VxBlock System analytics and manage capacity through integration with VMware vRealize Operations (vROps).

VxBlock Central Life Cycle Management enables inventory management for all Converged System components. This feature is integrated with CloudIQ, which enables you to monitor and analyze the inventories. Life Cycle Management provides the following information, which can help with system assessment and planning:

- A view of all the Converged Systems across data centers, from which you can analyze into component-level data
- An inventory of all the components in a Converged System organized by network, compute, storage, virtualization, and management
- Detailed software and firmware version information for all components in a Converged System
- Connections between components, where applicable

VxBlock Central options

VxBlock Central is available in the modular offerings of Base, VxBlock Central Workflow Automation, and VxBlock Central Advanced Analytics to manage your Converged Systems.

VxBlock Central Base option

The Base option enables you to monitor the health and compliance of VxBlock Systems through a central dashboard.

VxBlock System health monitors health or operational status of the following:

- The VxBlock System as a whole system
- The physical components such as a chassis, disk array enclosure, fan, storage processor, or X-Blade
- The compute, network, storage, and management components that logically group the physical components.

The landing page of VxBlock Central provides a view of the health and compliance of multiple VxBlock Systems. You can run a compliance scan on one or more VxBlock Systems. You can view key performance indicators for one or more components.

VxBlock Central contains dashboards that enable you to:

- View all components for selected VxBlock Systems, including detailed information such as serial numbers, IP address, firmware version, and location.
- View compliance scores and security and technical scan risks.

- View and compare RCMs on different systems.
- View real-time alerts for your system including severity, time, the system where the alert occurred, the ID, message, and status.
- Configure roles with AD integration.

The following table describes each dashboard:

Dashboard	Description
Inventory	<p>Provides the name, IP address, type, and element manager, RCM scan results, and alert count for components configured in a single view.</p> <p>An inventory item can be selected to suppress alerts and enable alerts. When alerts are suppressed for a specific component, real-time alert notifications are suspended.</p> <p>You can search VxBlock Systems for specific components or subcomponents and export a spreadsheet of your search.</p>
RCM	<p>Provides the compliance score, security, and technical risks associated with each VxBlock System. From the dashboard, you can:</p> <ul style="list-style-type: none"> • View noncompliant components, security, and technical risks associated for components. • Download software and firmware for your VxBlock System components to upgrade to a new RCM or remediate drift from your current RCM. • Run compliance scans and download and assess the results. • Check the base profile to determine whether components have the correct firmware versions. • Upload and install the latest compliance content. • Customize the compliance profile.
Alerts	<p>Provides real-time system alerts to monitor and receive alerts for critical failures on compute, storage, and network components. Version 2.0 and later provides VxBlock Central specific alerts for failures in system services such as discovery and Secure Remote Services. Administrators and Dell EMC Support can respond faster to incidents to minimize any impact of failures. Using the predefined alert notification templates to create alert notification profiles, you can specify how you want to be notified for a critical alert.</p>
Advanced Management	<p>Provides access to VxBlock Central Workflow Automation and VxBlock Central Advanced Analytics.</p>
Roles	<p>When VxBlock Central is integrated with AD, VxBlock Central authenticates AD users and supports mapping between AD groups and roles.</p> <p>Role mappings control the actions that a user is authorized to perform. By mapping an AD group to a role, you can control user permissions. When an AD user logs in to VxBlock Central, role mappings are checked for AD groups to which the user is assigned. The set of available permissions depends on the roles that are mapped to the groups in which the user is a member.</p>
Discovery	<p>Configure your VxBlock Systems in the MSM VM for monitoring and discovering of components with the Discovery feature. You can add manually or upload of <code>system.cfg</code>, and configure new Converged Systems and components, validate and discover components, and attach the VxBlock System to MSM VM, with the user interface.</p>

VxBlock Central Workflow Automation option

The Workflow Automation option provides automation and orchestration for daily provisioning tasks through integration with VMware vRealize Orchestrator (vRO).

Workflow Automation provides the following features:

- VxBlock Central Orchestration Services
- VxBlock Central Orchestration Workflows

The Workflow Automation option supports the following:

- VxBlock Systems 1000
- VxBlock System 740 and Vblock System 740
- VxBlock System 540 and Vblock System 540
- VxBlock System 350 and Vblock System 350

VxBlock Central Orchestration Services

VxBlock Central Orchestration Services manages credentials and uploads ISO images for server installation. VxBlock Central Orchestration Services manages the following services:

- Common operations
- Credential store
- Firmware repo
- Gateway
- VMware Service Manager operations
- PXE configuration
- Registry
- Storage operations

The VxBlock Central Orchestration VMware vRO Adapter provides supported workflows for VxBlock System compute expansion.

VxBlock Central Orchestration Workflows

VxBlock Central Orchestration Workflows simplify complex compute, storage, and network provisioning tasks using automated workflows for VMware vRO.

Automated VxBlock Central Orchestration Workflows enable you to concurrently provision multiple VMware vSphere ESXi hosts and add these hosts to the VMware vCenter cluster. The workflow provides the validation and resilience that is required for enterprise-grade operations. Once hosts are provisioned, workflows trigger an RCM compliance scan to ensure compliance with RCM standards. The VMware vRO workflows also support bare-metal server provisioning.

The following table provides information about the different types of workflows available through VMware vRO:

Workflow	Description
Configuration	Establishes the connection between VMware vRO with automation workflows and VxBlock Central Orchestration Services to run workflow automation.
Service	Provides a presentation layer for user input and data validation. Service workflows create multiple instances of fulfillment workflows to run concurrently.
Fulfillment	After you run a service workflow, a corresponding fulfillment workflow is created in the Fulfillment directory of the VxBlock System. If you experience a failure, the fulfillment workflow enables you to troubleshoot the cause and fix the issue.

VxBlock Central Advanced Analytics option

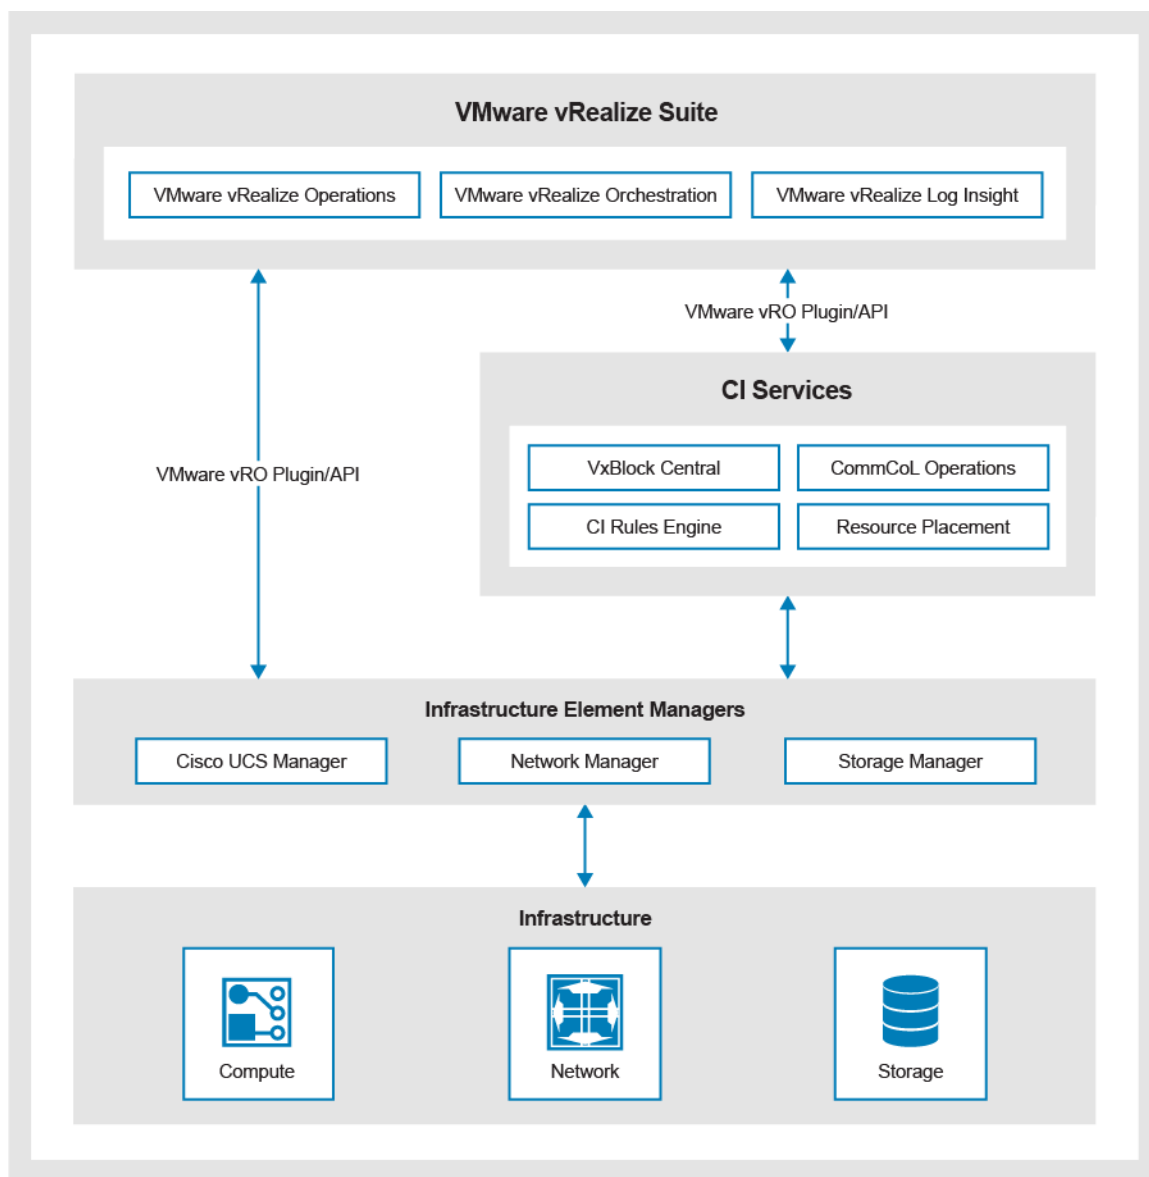
The Advanced Analytics option provides features that simplify operations you must perform for VxBlock Systems through advanced monitoring, system analytics, and simplified capacity management.

VMware vRealize Operations (vROps) Manager integration with VxBlock Central presents the topology and relationship of VxBlock Systems with compute, storage, network, virtualization, and management components. Advanced Analytics provides advanced monitoring, system analytics, and simplified capacity management through integration with VMware vROps Manager.

Advanced Analytics enables you to:

- Monitor health, performance, and capacity through predicative analytics.
- Troubleshoot and optimize your environment through alerts and recommended actions.
- Manage inventory and create reports.
- Define custom alerts for performance and capacity metrics in the following actions:
 - Collect data from VxBlock Systems every 15 minutes by default.
 - Collect real-time alerts from VxBlock Systems every three minutes, by default.
 - View VxBlock Central VM relationships to physical infrastructure. Core VM, MSM VM, and MSP VM resource monitoring enables you to identify and monitor a collection of resources that are associated with a VM.

The following figure provides an overview of how VxBlock Central uses VMware vRealize:



VxBlock Central architecture

VxBlock Central uses VMs to provide services.

The following table provides an overview of VxBlock Central VMs:

VM	Description
Core	Discovers and gathers information about the inventory, location, and health of the VxBlock System
MSM	Provides functions to manage multiple VxBlock Systems In a data center environment, one MSM VM can be associated with up to 8 Core VMs
MSP (optional)	Provides functions for RCM content repositioning.

VxBlock Central includes the Core VM and the multisystem management (MSM) VM as a minimum configuration. The multisystem repositioning (MSP) VM deployment is optional for repositioning.

Discovery

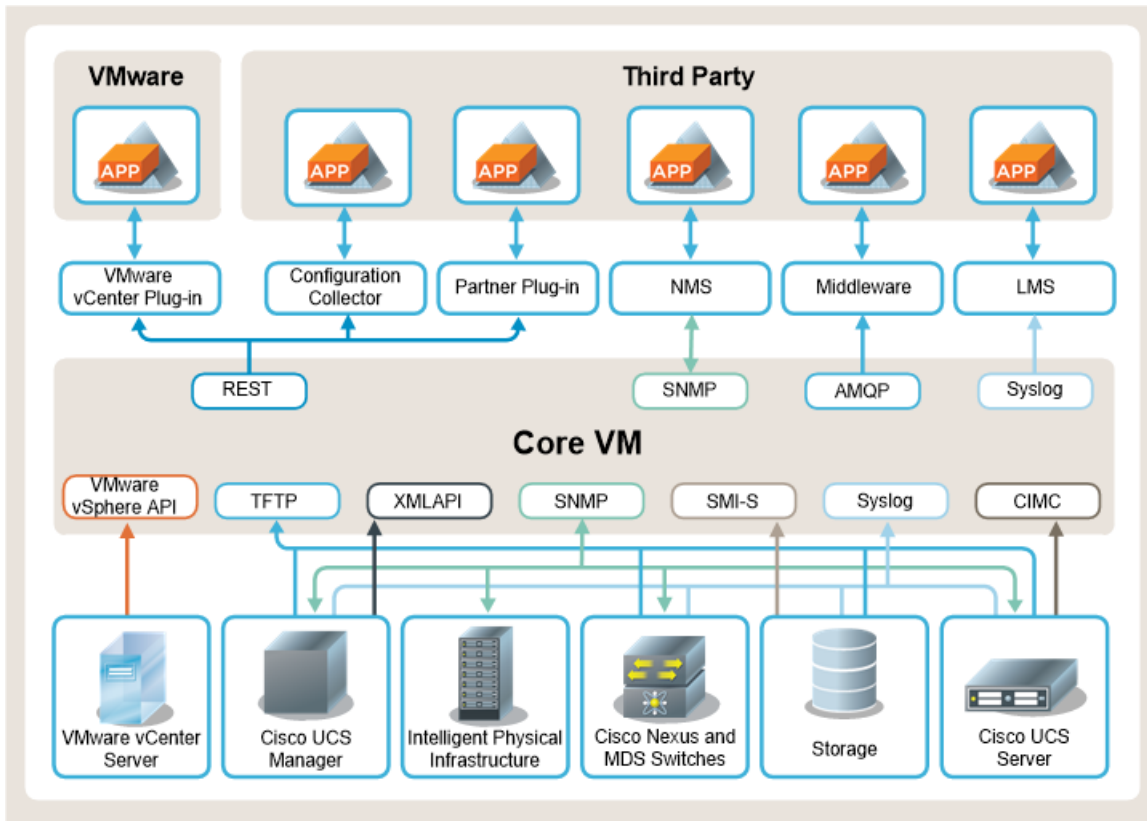
The discovery model resides in a database and is exposed through REST and SNMP interfaces. Initial discovery is performed during manufacturing of the VxBlock System and relies on an .XML file that contains build and configuration information. Core VM uses the .XML file to populate basic information about the VxBlock System and establish communication with components.

After initial discovery, Core VM uses the following methods to discover the VxBlock System, including physical components and logical entities:

- XML API
- SNMP
- SMI-S
- Vendor CLIs, such as Unisphere CLI
- Platform Management Interface

Core VM performs discovery every 15 minutes, by default. This setting can be changed.

The following figure is a high-level overview of integration between Core VM and various products and protocols:



Data collection

VxBlock Central uses data collectors to extract required data from various web services.

The following table describes the data collectors:

Data collector	Description
VxBlock Central collector	Uses the VxBlock Central REST API to collect the VxBlock System configuration data and key performance indicators (KPI) already discovered in Core VM. The configuration is stored with KPI data from Core VM into the Cassandra and Elasticsearch databases.
SMI-S collector	Works with the CIM Object Manager (ECOM) service that runs on SMI components to discover metrics for VMAX: <ul style="list-style-type: none"> • Storage array • Storage processor • Storage volume • Storage pool • Storage tier • Disk
SNMP collector	Collects information from SNMP enabled devices such as Cisco Nexus and MDS switches to discover metrics. Information can be collected from the following network components: <ul style="list-style-type: none"> • Switches • Network chassis • Container • Fan • Expansion module • Power supply bay • PSU • Network temperature sensor • SFP • IPI appliance
vSphere API collector	Works with VMware vCenter Server using the VMware vSphere API to discover metrics, for datastores, disk partitions, and clusters.
Dell EMC Unity REST collector	Collects configuration data from Dell EMC Unity and Unity XT storage array and its components.
XIO REST collector	Collects metrics for storage array, storage volume, disk, and port. VxBlock Central collects all other configuration information with the collector.
XML API collector	Collects information from the Cisco UCS using the XML API to discover metrics.
VMware NSX collector	Collects information about VMware NSX components, such as Virtual Appliance Management and the NSX controllers. The NSX collector interfaces with the NSX Manager APIs.

VxBlock Central Shell

The VxBlock Central Shell removes the complexity of working with individual component interfaces and provides a plug-in structure that can be extended to include more functionality. VxBlock Central Shell creates an abstraction layer that removes the burden of having to use different login credentials, IP addresses, and syntax to make configuration changes across multiple components. VxBlock Central Shell can help manage multiple VxBlock Systems.

For example, to update the NTP server IP addresses for all switches on a VxBlock System, you can issue a single command without having to log in to each component.

```
>%ntp switch set ['10.1.139.235', '10.1.219.13']
[Switch 'N5B' at 10.1.139.23:, result: ['10.1.139.235', '10.1.219.13'],
Switch 'N5A' at 10.1.139.22:, result: ['10.1.139.235', '10.1.219.13'],
Switch 'MGMT-N3B' at 10.1.139.2:, result: ['10.1.139.235', '10.1.219.13'],
Switch 'MGMT-N3A' at 10.1.139.1:, result: ['10.1.139.235', '10.1.219.13'],
Switch 'N1A' at 10.1.140.235:, result: ['10.1.139.235', '10.1.219.13'],
Switch 'M9A' at 10.1.139.20:, result: ['10.1.139.235', '10.1.219.13'],
Switch 'M9B' at 10.1.139.21:, result: ['10.1.139.235', '10.1.219.13']]
```

The shell is a framework layer built on top of Python and VxBlock Central API bindings. In addition to the commands provided, any valid Python command can be run in the shell.

Developers writing extensions for the VxBlock Central Shell can provide a single interface for all components and enable users to:

- Perform operations on each VxBlock System as a single logical entity rather than a collection of components.
- Configure and manage settings at the individual VxBlock System component level.

Secure Remote Services

VxBlock Central can connect to Secure Remote Services (SRS) to automatically send system inventory, real-time alerts, and RCM fitness information through the SRS connection to collect and analyze data.

Note: Customers using VxBlock Central Version 2.0 and later configure SRS in the VxBlock Central end-user interface. For more information, refer to Configure SRS gateway in *VxBlock Central Version 2.0 Online Help*.

Customers on VxBlock Central Version 1.5 and earlier use VxBlock Central Shell SRS Extension Pack to configure SRS.

Use the appropriate procedure to perform the following functions:

- Configure VxBlock Central to use SRS.
- Deregister VxBlock Central with SRS.
- Update a SRS gateway configuration or VxBlock Central ID (SWID).
- Uploads the following information to SRS about your Converged System:
 - Release Certification Matrix (RCM) compliance scan results (ZIP file containing CSV, XLS, PDF, and XML files) (if you have installed RCM content and selected a default profile).
 - Converged System inventory files (JSON).
 - Converged System real-time alerts are automatically sent to SRS if notification is configured.
- Modify the schedule VxBlock Central uses to regularly send RCM and inventory information to SRS (only in VxBlock Central Version 1.5 and earlier).
- Edit settings to send data to SRS (only in VxBlock Central Versions 2.0 and later).

Key performance indicators

Access key performance indicator (KPI) information using VxBlock Central or MSM VM. VxBlock Central displays charts and graphs of KPI information for the selected element type.

The following table provides examples of KPI information:

Element type	Examples of KPI information
storagearray	Remaining raw capacity
	Total space available for user data
	Remaining managed space

Element type	Examples of KPI information
	Total IO per second
storagepool	User capacity
disk	Disk raw capacity
	Bandwidth
switch	Current bandwidth
	Number of error inbound packets
rack	Monitor total energy.
	Monitor average power consumption.
computeserver	Total memory
	Temperature

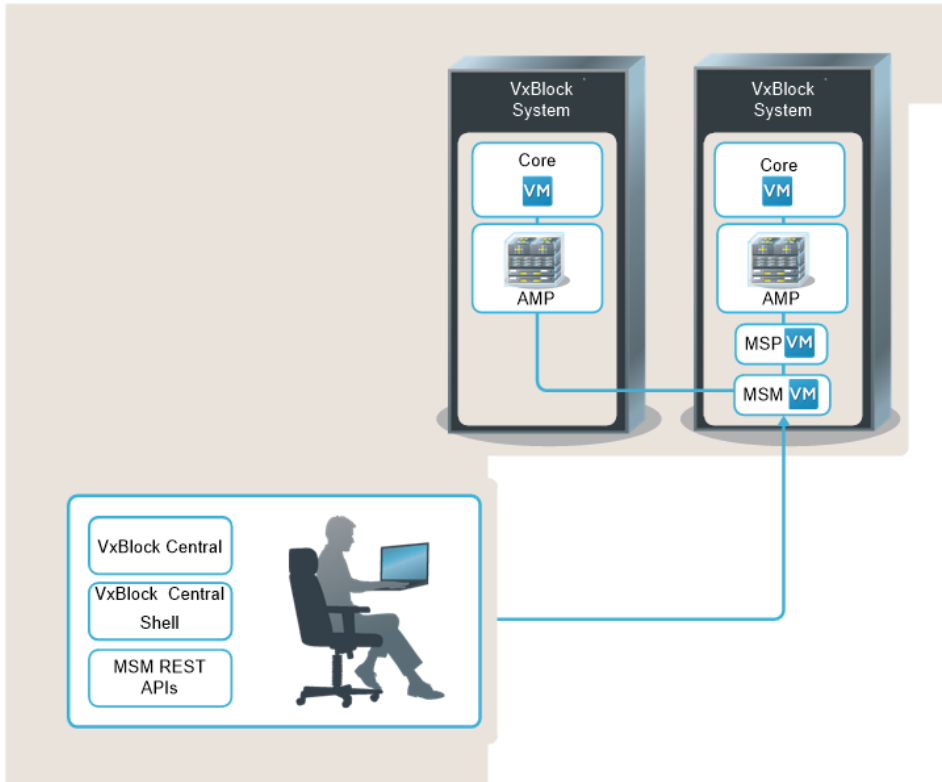
The MSM VM API for multisystem services retrieves the following KPI data:

- All existing KPI definitions in the VxBlock System.
- Existing KPI definitions for a particular element type and/or component category.
- Time series KPI data:
 - A particular time resolution.
 - A start time for time series queries.
 - An end time for time series queries.

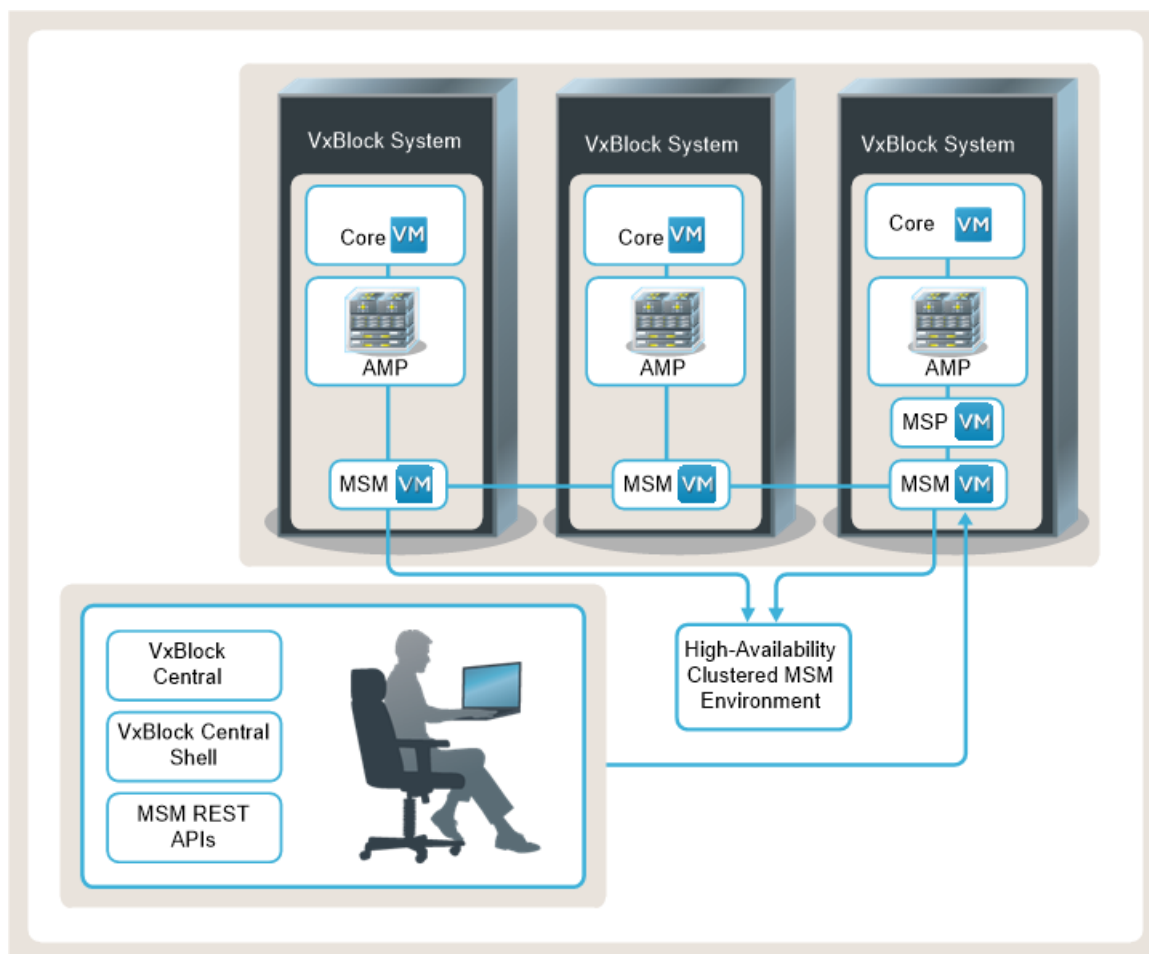
Data center architecture

VxBlock Central supports a clustered environment that includes multiple MSM VMs configured to run in a single physical data center or in multiple, geographically separate physical data centers.

In a data center environment, one MSM VM can be associated with up to eight Core VMs:



The following figure shows a single-site environment consisting of three MSM VMs, each associated with a single Core VM:



MSM VMs are configured to form a cluster. Capabilities and functionality are exposed after deployment and configuration.

In a single-site configuration with one datacenter, VxBlock Central supports up to three MSM VMs running in the data center. Up to eight Core VMs are supported.

VxBlock Central supports a multisite clustering configuration that includes a maximum of three data centers. Up to two MSM VMs are supported.

AMP overview

VxBlock Systems use AMP Central, AMP-VX or AMP-3S to manage local and production workloads without exhausting system resources.

AMP provides the ability for VxBlock System to:

- Manage a single VxBlock System or multiple VxBlock Systems.
- Run the core and Dell EMC optional management workloads.
- Monitor and manage health, performance, and capacity.
- Provide network and fault isolation for management.
- Eliminate resource overhead.

See the *Dell EMC AMP Central Production Guide* for more information about AMP Central.

The following table provides an overview of VxBlock Central support for the AMPs:

AMP	VxBlock Central Base	VxBlock Central Workflow Automation	VxBlock Central Advanced Analytics
AMP-3S	Yes	Yes	Yes
AMP-VX	Yes	Yes	Yes
AMP Central	Yes	Yes	Yes

AMP hardware components (VMware vSphere 6.x)

AMP systems support management software that help install and operate components within VxBlock Systems 1000.

The following table provides an overview of the AMP hardware components:

AMP	Physical server	Storage	Description
AMP-3S	For VMware vSphere 6.7: <ul style="list-style-type: none"> 2 to 16 Cisco UCS C220 M4 (3 required for HA)** 2 to 16 Cisco UCS C220 M5 servers (3 required for HA)* For VMware vSphere 6.5: <ul style="list-style-type: none"> 2 to 6 Cisco UCS C220 M4 servers 2 to 16 Cisco UCS C220 M5 servers* 	<ul style="list-style-type: none"> FlexFlash SD for VMware vSphere ESXi boot Dell EMC Unity with FAST Cache and FAST VP 	Provides limited scalability configuration using Cisco UCS C220 servers and extra storage expansion capacity.
AMP-VX	4 to 8 PowerEdge R640 servers	<ul style="list-style-type: none"> SD card for VMware vSphere ESXi boot VMware vSAN all-flash shared storage 	AMP-VX configuration of 4 servers supports the core management workload for the AMP-VX and two VxBlock System. Servers can be added to manage multiple VxBlock Systems.
AMP Central	4 to 16 Cisco UCS C220 M5 servers	<ul style="list-style-type: none"> FlexFlash SD for VMware vSphere ESXi boot Dell EMC Unity with FAST Cache and FAST VP 	Provides limited scalability configuration using Cisco UCS C220 M5 servers and extra storage expansion capacity.

*AMP-3S M4/M5 supports the VxBlock Central Base option with two servers. VxBlock Central Workflow Automation and VxBlock Central Advanced Analytics options require a minimum of three AMP-3S servers.

**AMP-3S deployed using Cisco UCS C220 M4 servers can be expanded with Cisco UCS C220 M5 servers using VMware EVC.

AMP software components (VMware vSphere 6.x)

AMPs are delivered with specific installed software components dependent on the selected Release Certification Matrix (RCM).

The following software components are installed dependent on the selected RCM:

- Microsoft Windows Server

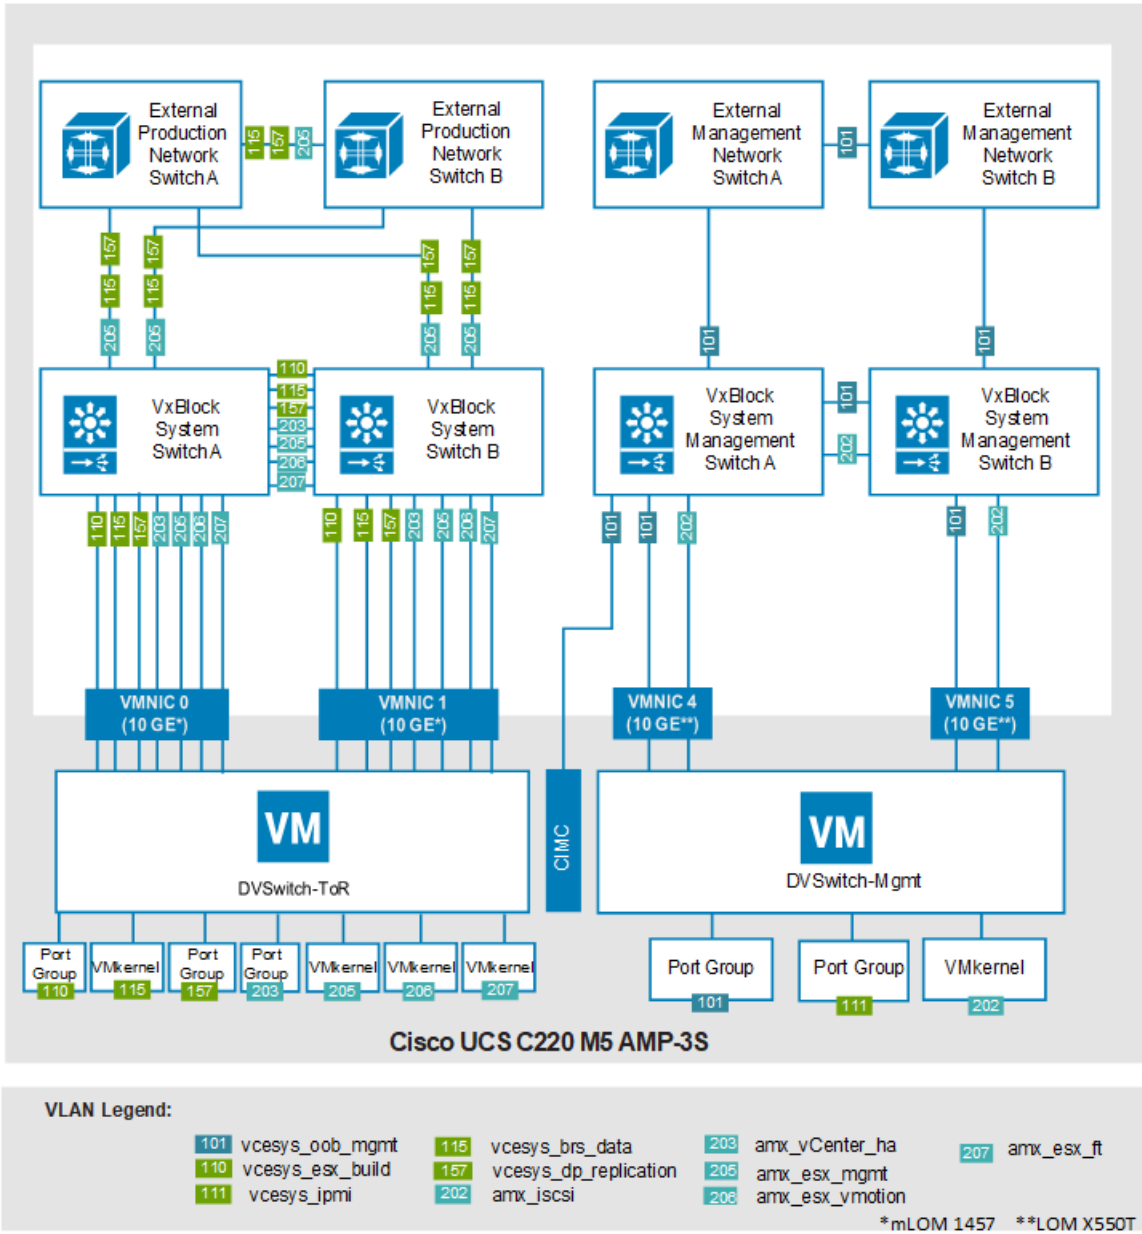
- VMware vSphere Enterprise Plus
- VMware vSphere Hypervisor ESXi
- VMware Single Sign-On (SSO) Service
- Optional: VMware vSphere Platform Services Controller (PSC)
- VMware vSphere Web Client Service
- VMware vSphere Inventory Service
- VMware vCenter Server Appliance (vCSA)
- **Note:** For VMware vSphere 6.x, only the VMware vCSA deployment model is offered.
- VMware vCenter Update Manager (VUM – Integrated with VMware vCSA)
- VMware vSphere client
- VMware vSphere Web Client (Flash/Flex client)
- VMware vSphere Syslog Service (optional)
- VMware vSphere Core Dump Service (optional)
- VMware vSphere Distributed Switch (VDS)
- PowerPath/VE Management Appliance (PPMA)
- Secure Remote Services (SRS)
- Array management modules, including but not limited to:
 - Unisphere Client
 - Unisphere Service Manager
 - Dell EMC Unity Initialization Utility
 - SMI-S Provider
- AMP-VX and AMP-3S: Cisco Prime Data Center Network Manager (DCNM) and Device Manager
- (Optional) RecoverPoint management software that includes the management application and deployment manager
- Embedded VMware PSCs are supported with VMware vSphere 6.7 and 6.5
- VMware vCenter HA is supported with VMware vSphere 6.7 and 6.5

Management network connectivity (VMware vSphere 6.x)

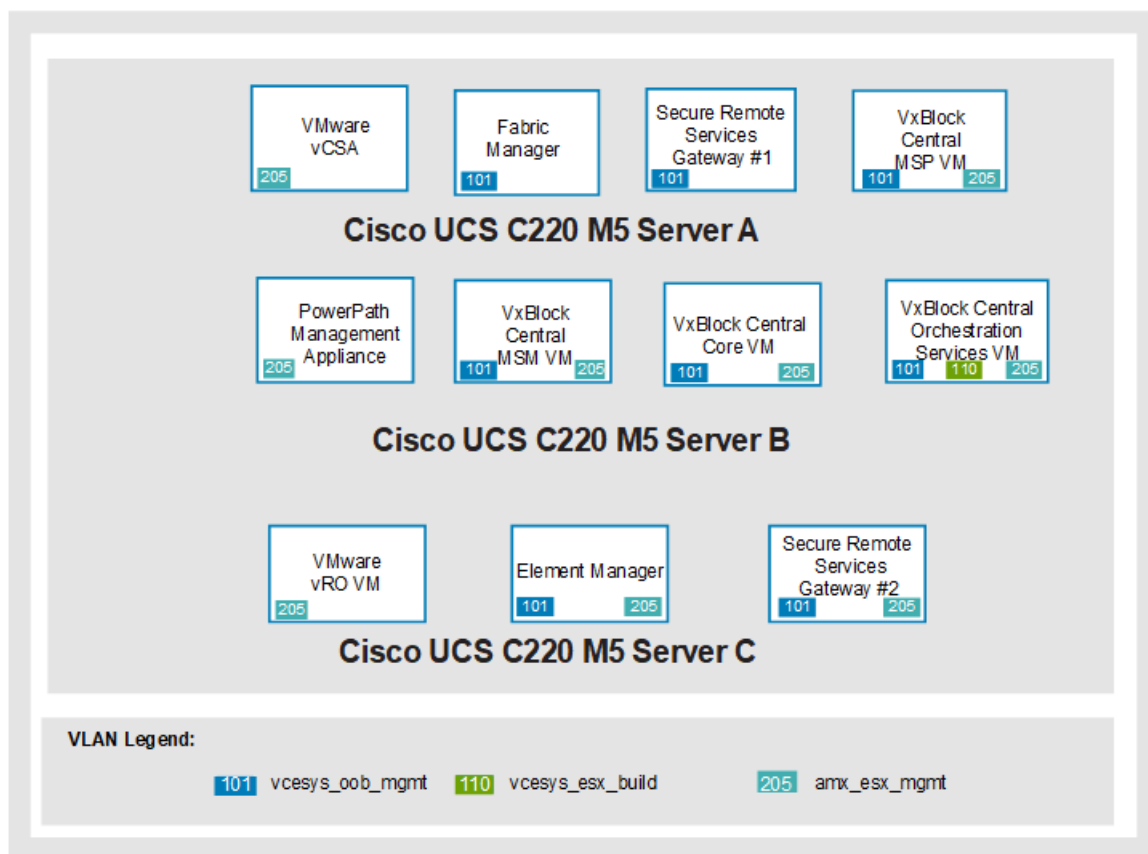
Network connectivity and server assignments on Cisco UCS C220 M5 servers are provided for VxBlock Systems with VMware vSphere 6.7 and 6.5.

AMP-3S

The following figure provides network connectivity on Cisco UCS C220 M5 servers with VMware vSphere 6.7:



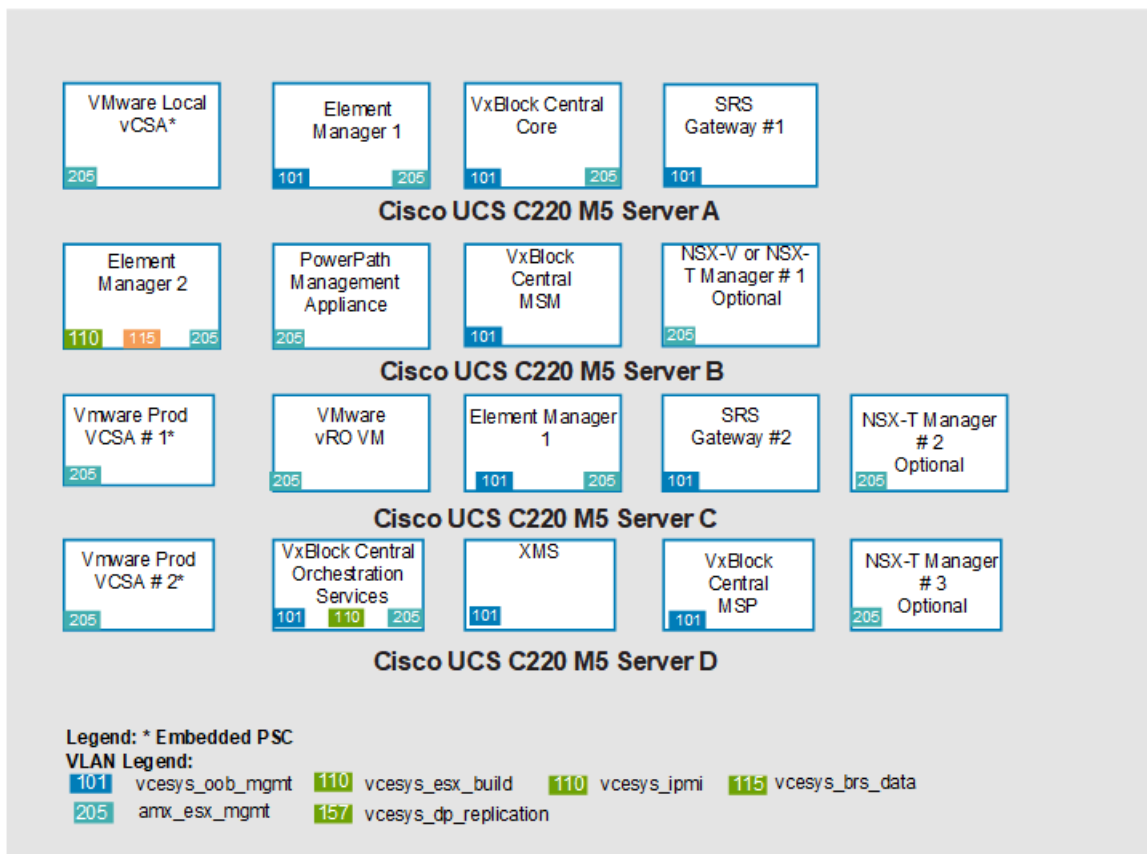
The following figure shows VM placement with three AMP-3S servers:



VLANS 101, 205, and 115 support an extra Element Manager VM for Integrated Data Protection. Deploy both Element Manager VMs on different C220 M5 servers.

AMP Central

AMP Central uses VMware VDS with the network input output configuration (NIOC) in place of VMware vSphere standard switches. The following figure shows the VM placement for four Cisco UCS C220 M5 servers:



VLANs 101, 205 and 115 support an extra element manager VM for data protection. Deploy Element Manager VMs on different Cisco UCS C220 M5 servers.

Virtualization layer (VMware vSphere 6.x)

VMware vSphere 6.5 or 6.7 is the virtualization platform that provides the foundation for the private cloud.

Virtualization components (VMware vSphere 6.x)

VMware vSphere ESXi and VMware vCenter Server are the core components for virtualization management.

VMware vSphere 6.x includes a pair of VMware Platform Service Controller (PSC) Linux appliances to provide the VMware Single Sign-on (SSO). With VMware vSphere 6.7, VMware vCenter Server Appliance (vCSA) is the default deployment model for VMware vCenter Server. With VMware vSphere 6.5, VMware vCSA with embedded VMware PSC is the default deployment model for VMware vCenter Server.

Note: VMware vSphere 6.7 and VMware vSphere 6.5 are the last VMware vSphere releases to support the deployment of external VMware PSCs. Updates to both of these VMware vSphere versions support Enhanced Linked Mode (ELM) in embedded PSCs, making external VMware PSCs no longer necessary. Upgrading to VMware vSphere 6.7 update 1 or VMware vSphere 6.5 update 2d enables you to migrate external VMware PSCs to embedded PSCs. For more information about changing to embedded VMware PSCs from external VMware PSCs, consult your account team.

The hypervisors are deployed in a cluster configuration. The cluster enables dynamic allocation of resources, such as CPU, memory, and storage. The cluster also provides workload mobility and flexibility with the use of VMware vSphere vMotion and Storage vMotion technology.

VMware vSphere Hypervisor ESXi (VMware vSphere 6.x)

The VMware vSphere Hypervisor ESXi runs on the Management Appliance and VxBlock Systems servers using VMware vSphere Enterprise Plus.

The lightweight hypervisor requires little space to run (less than 6 GB of storage required to install) with minimal management overhead.

By default, the hypervisor is installed on a 64 GB Cisco FlexFlash SD Card (mirrored HV partition). All Cisco FlexFlash (boot) capable hosts are configured with a minimum of two 64 GB SD cards.

The compute hypervisor supports 4–6 10 GbE physical NICs (pNICs) on the VxBlock System 1000 VICs.

VMware vSphere ESXi does not contain a console operating system. The VMware vSphere Hypervisor ESXi boots from Cisco FlexFlash SD cards. The stateless hypervisor (PXE boot into memory) is not supported.

Cluster configurations

VMware vSphere ESXi hosts and their resources are pooled together into clusters.

These clusters contain the CPU, memory, network, and storage resources available for allocation to VMs. Clusters can scale up to a maximum of 64 hosts, enabling the support for thousands of VMs.

The clusters can also support various Cisco UCS blades running inside the same cluster. Some advanced CPU functionality might be unavailable if more than one blade model is running in a given cluster.

Datstores

The VxBlock System 1000 supports a mixture of datastore types: block level storage using VMFS or file level storage using NFS.

The maximum size per VMFS5/VMFS6 volume is 64 TB (50 TB VMFS3 @ 1 MB), with a maximum VMDK file size of 62 TB. Each host can support a maximum of 512 VMFS volumes and/or a maximum of 256 NFS datstores.

Dell EMC optimizes advanced settings for VMware vSphere ESXi hosts deployed in the VxBlock 1000 to maximize the throughput and scalability of VMFS and NFS datstores.

Virtual networks

VMware Virtual Distributed Switch (VDS) provides virtual networking and uses a minimum of four uplinks presented to the hypervisor.

vNICs are equally distributed across all available physical adapter ports to ensure redundancy and maximum bandwidth where appropriate. This distribution provides general consistency and balance across all Cisco UCS B-series and C-series models, regardless of the Cisco UCS Virtual Interface Card (VIC) hardware.

VMware vSphere ESXi has a predictable uplink interface count. All applicable VLANs, native VLANs, MTU settings, and QoS policies are assigned to the vNIC to ensure consistency if uplinks require migration to VMware VDS after manufacturing.

VMware vCenter Server (VMware vSphere 6.x)

VMware vCenter Server 6.x resides on the VMware vCenter Server Appliance (vCSA) and manages hypervisors and VMs.

VMware vCenter Server is deployed by default on the VMware vCSA. VMware vCenter Update Manager is fully integrated with the VMware vCSA and runs as a service to assist with host patch management.

In an external deployment for VMware vSphere 6.7 or for a regular deployment of VMware vSphere 6.5, AMP Central or AMP-VX and the VxBlock System 1000 each have a VMware vCSA instance with a pair of VMware Platform Service Controllers (PSCs). For VMware vSphere 6.5, VMware PSCs are optional.

In an embedded VMware PSC deployment for VMware vSphere 6.7, AMP Central or AMP-VX and the VxBlock System 1000 have only the VMware vCSA.

The AMP Central or AMP-VX management appliance hosts the VMware vCenter Server and the VxBlock System VMware vCenter instances.

The VMware vCenter Server provides the following functionality for AMP-VX, AMP Central, or AMP-3S:

- Clone VMs
- Create templates
- Monitor hosts and VMs
- Creates alerts
- VMware vSphere vMotion and VMware Storage vMotion
- Configuration of VMware Distributed Resource Scheduler (DRS) and VMware vSphere HA clusters

AMP-3S supports a single VMware vCSA to manage the local AMP resources and the platform workload resources. AMP Central has a dedicated VMware vCSA to manage local AMP resources and has an extra VMware vCSA or VMware vCSAs to manage platform workload resources.

VxBlock System administrators can create and apply the following alarms to all managed objects in VMware vCenter Server:

- Data center, cluster, and host health, inventory, and performance

- Data store health and capacity
- VM usage, performance, and health
- Virtual network usage and health

Databases

The VMware vCSA uses the embedded PostgreSQL database. The VMware vCenter Update Manager (VUM) and the VMware vCSA share the same PostgreSQL database server, but use separate PostgreSQL database instances.

Authentication

The VxBlock System supports the VMware Single Sign-On (SSO) capable of the integrating the following multiple identity sources:

- AD
- Open LDAP
- Local accounts for authentication

VMware vSphere 6.7 includes a pair of VMware PSC Linux appliances to provide the SSO service. In an embedded deployment, the VMware PSC is integrated into the VMware vCSA. VMware vCenter Server, Inventory, Web Client, SSO, Core Dump Collector, and Update Manager run as separate services. Each service can be configured to use a dedicated service account depending on the security and directory services requirements.

VMware vCenter Server features

Dell EMC supports the following VMware vCenter Server features:


- VMware SSO
- VMware PSC
- VMware vSphere Web Client (used with VxBlock Central)
- VMware vSphere Distributed Switch (VDS)
- VMware vSphere HA
- VMware DRS
- VMware Fault Tolerance
- VMware vSphere vMotion
- VMware Storage vMotion
- VMware Storage vMotion (Layer 3 capability available for compute resources)
- Raw Device Mappings
- Resource Pools
- Storage DRS (capacity only)
- Storage driven profiles (user-defined only)
- Distributed power management (up to 50 percent of VMware vSphere ESXi hosts/blades)
- VMware Syslog Service
- VMware Core Dump Collector
- VMware vCenter Web Client

VMware vSphere 6.x enables support of the following other VMware vCenter Server features:

- Embedded deployment
- Embedded linked mode
- VMware vCenter high availability simple deployment

Virtualization layer (VMware vSphere 6.x)

- VMware vCenter high availability advanced deployment
- VMware vSphere HTML5 Client for VMware vSphere 6.7 only

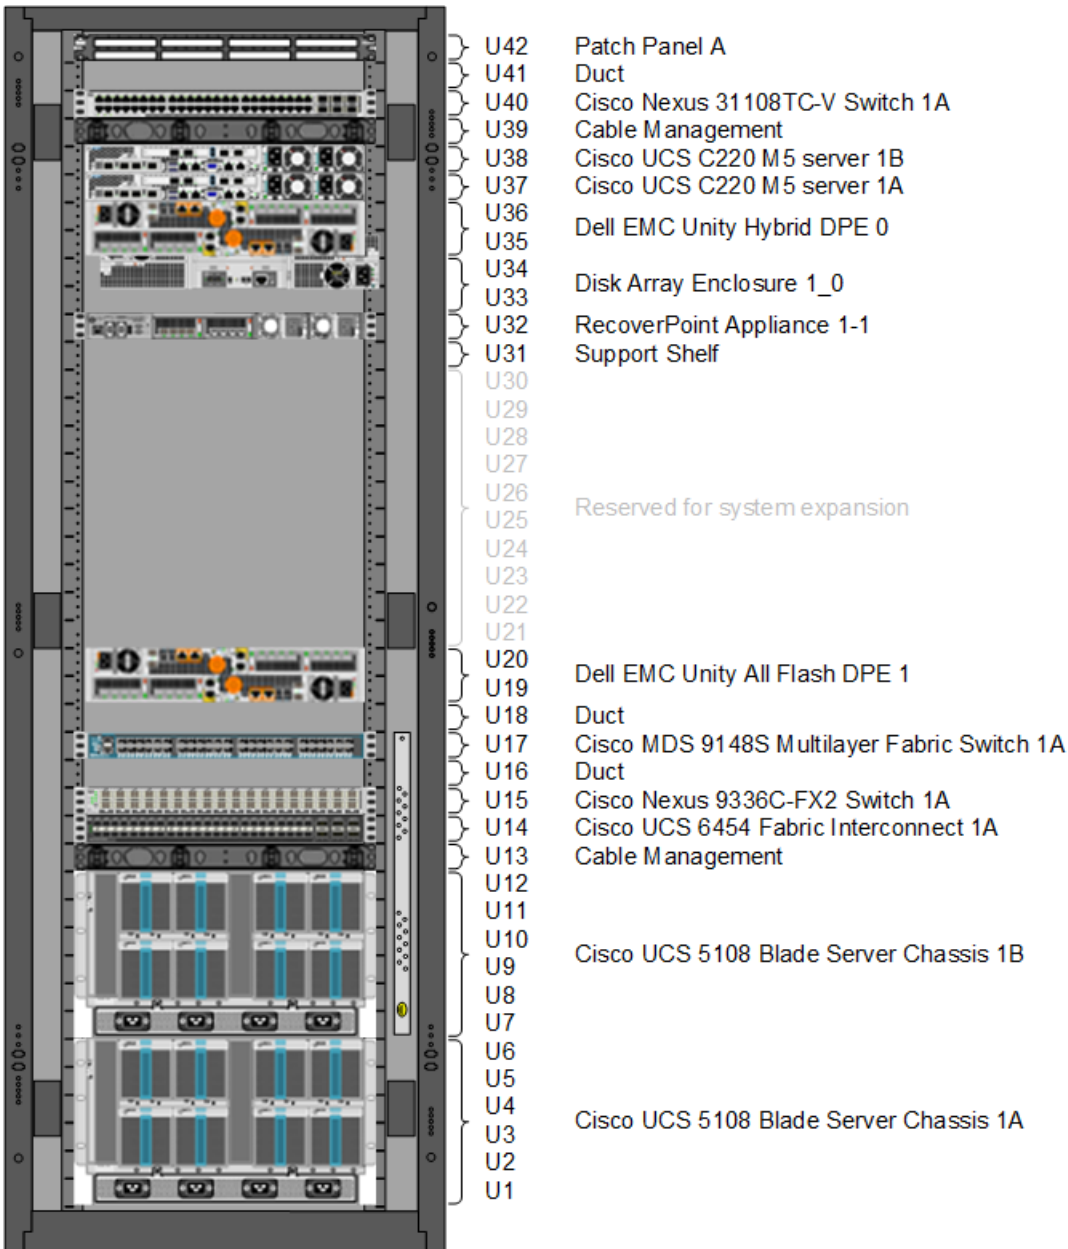
 **Note:** All engineered guidance and procedures for VMware vCenter high availability assume four hosts.

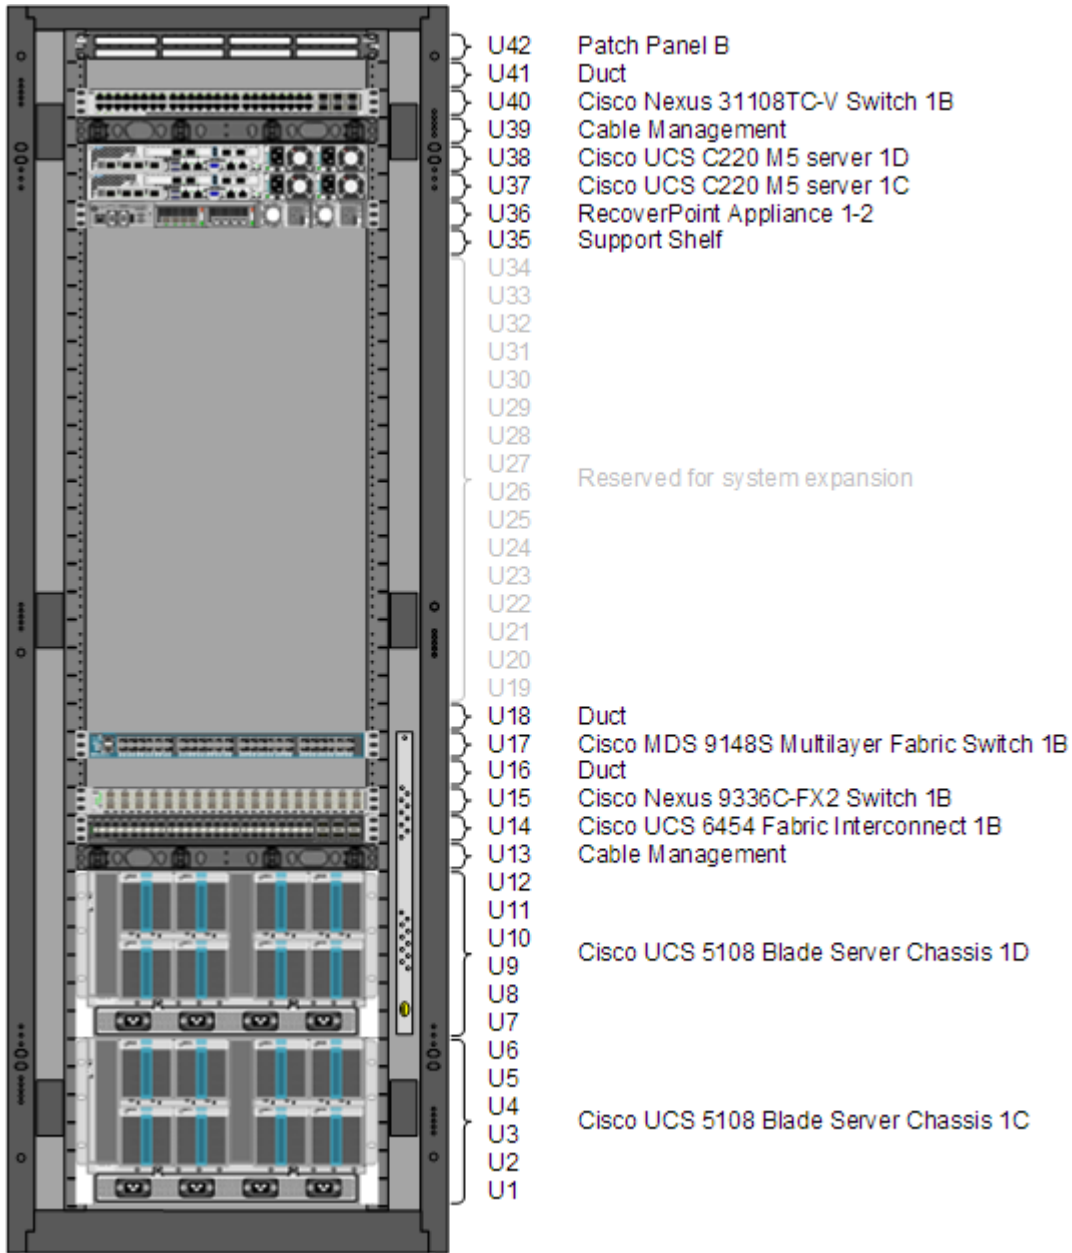
VxBlock System 1000 sample configurations

Cabinet elevations vary based on the specific configuration requirements.

VxBlock System 1000 with AMP Central integrated configuration

Elevations are provided for sample purposes only. For specifications for a specific design, consult your Dell EMC Sales Engineer.

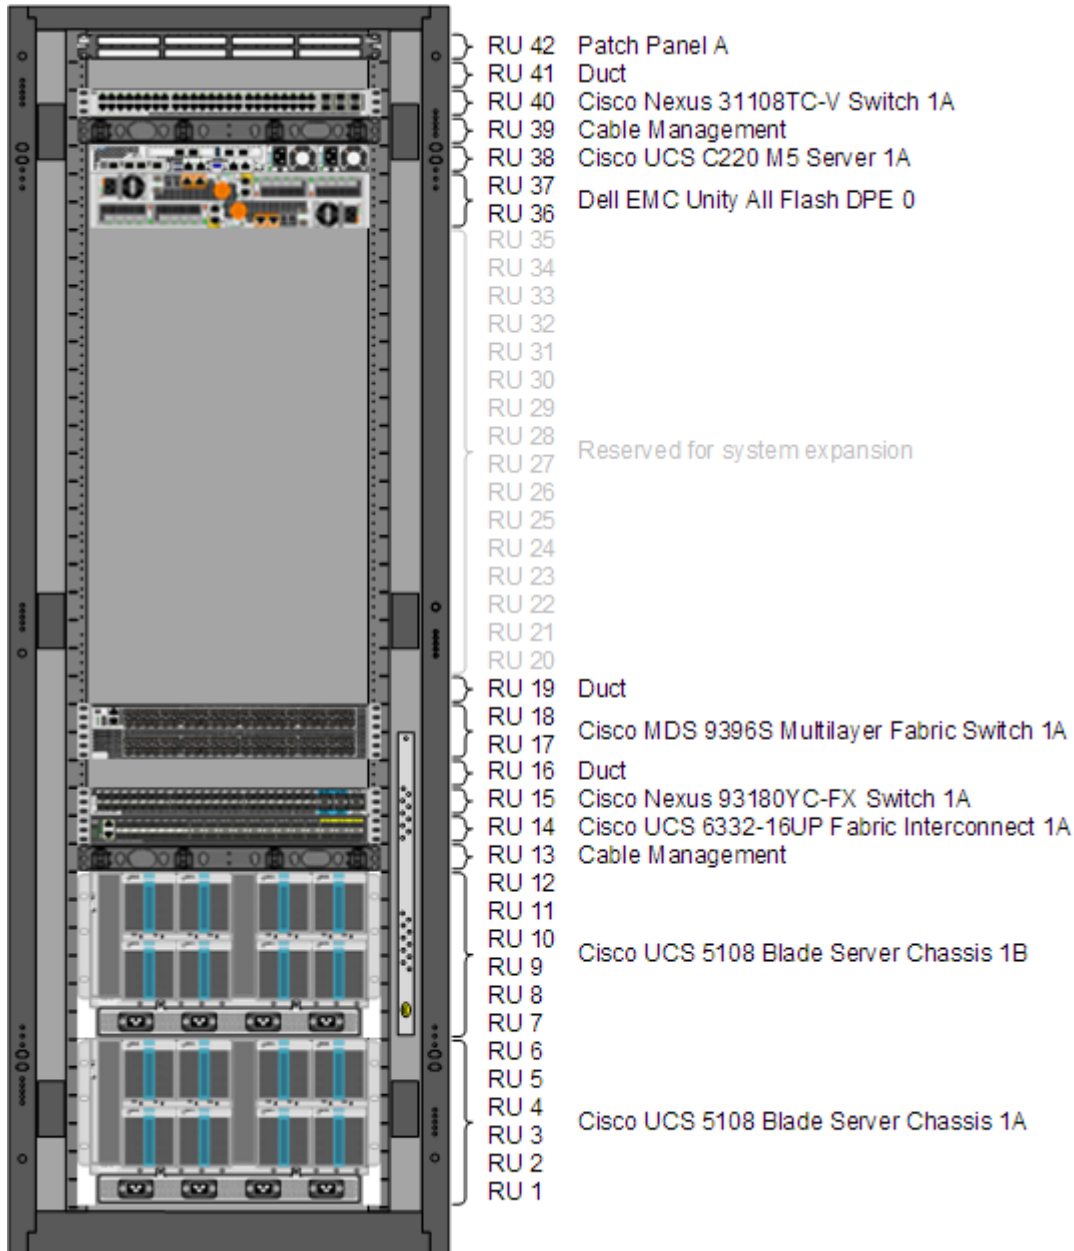




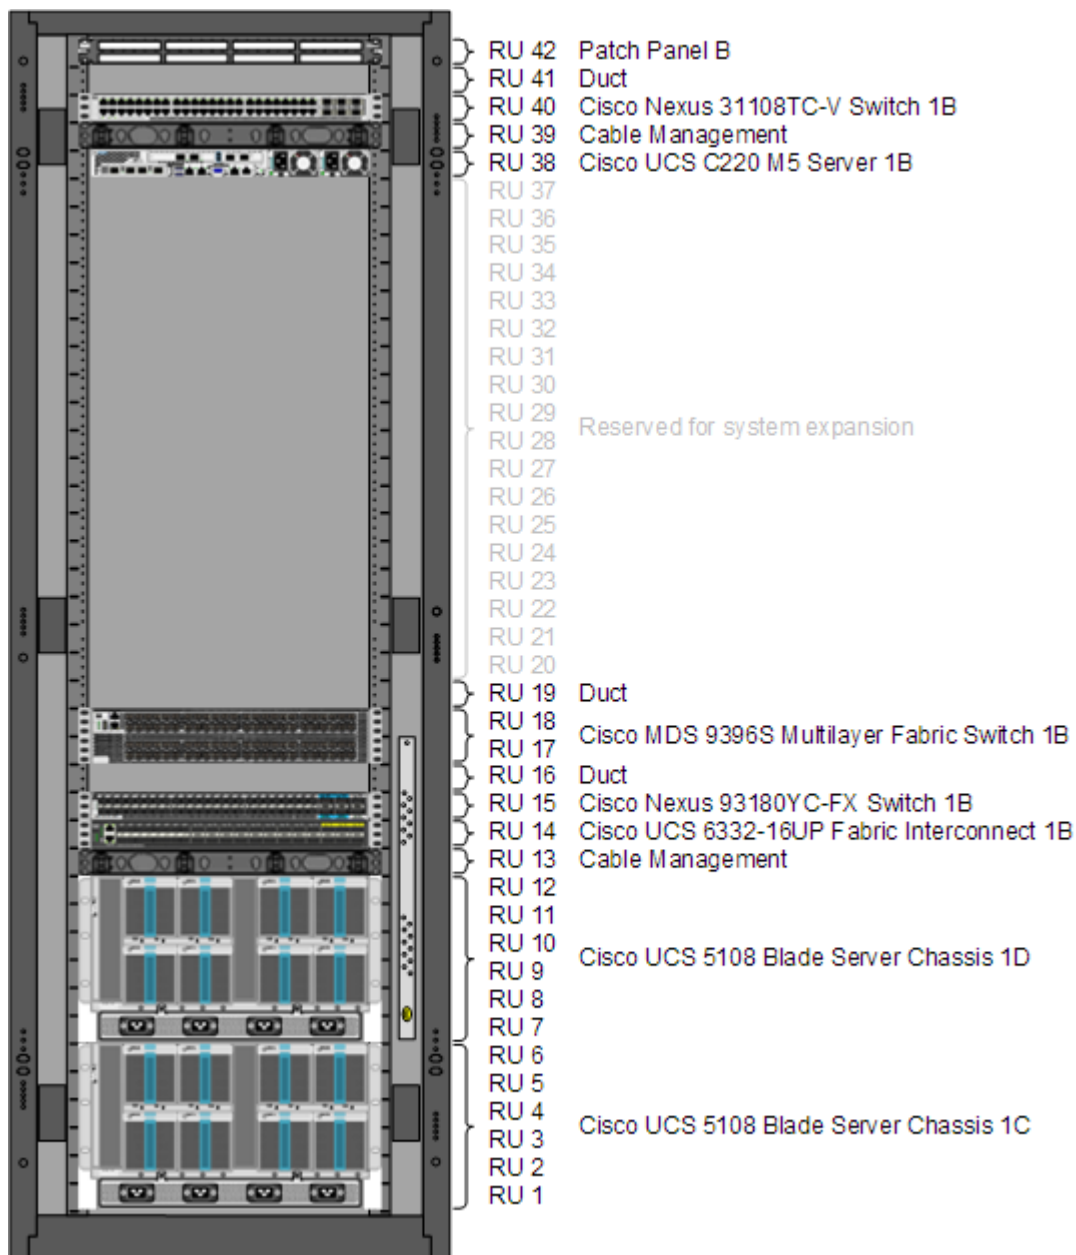
VxBlock System 1000 with AMP-3S

Elevations are provided for sample purposes only. For specifications for a specific design, consult your Dell EMC Sales Engineer.


Cabinet 1



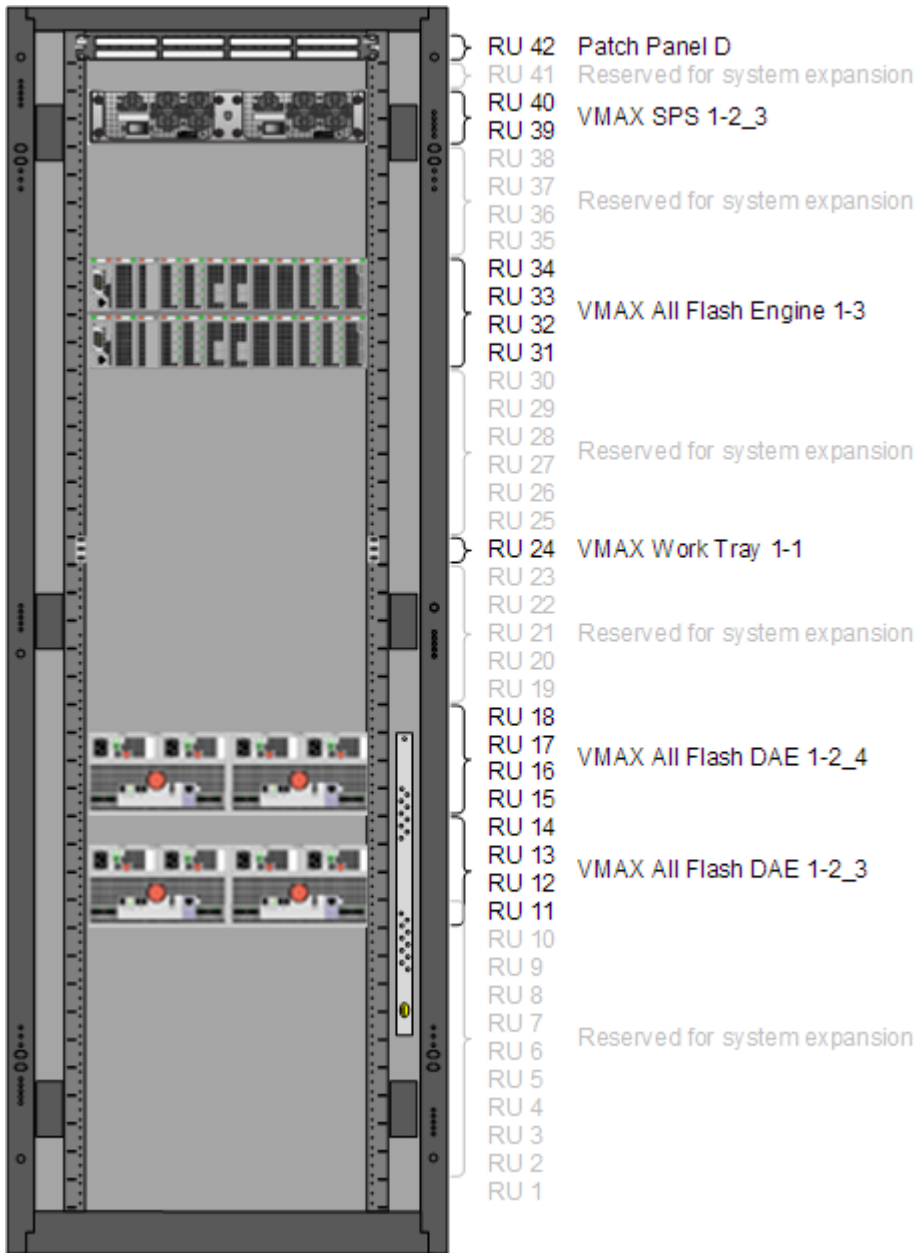
Cabinet 2



Cabinet 3

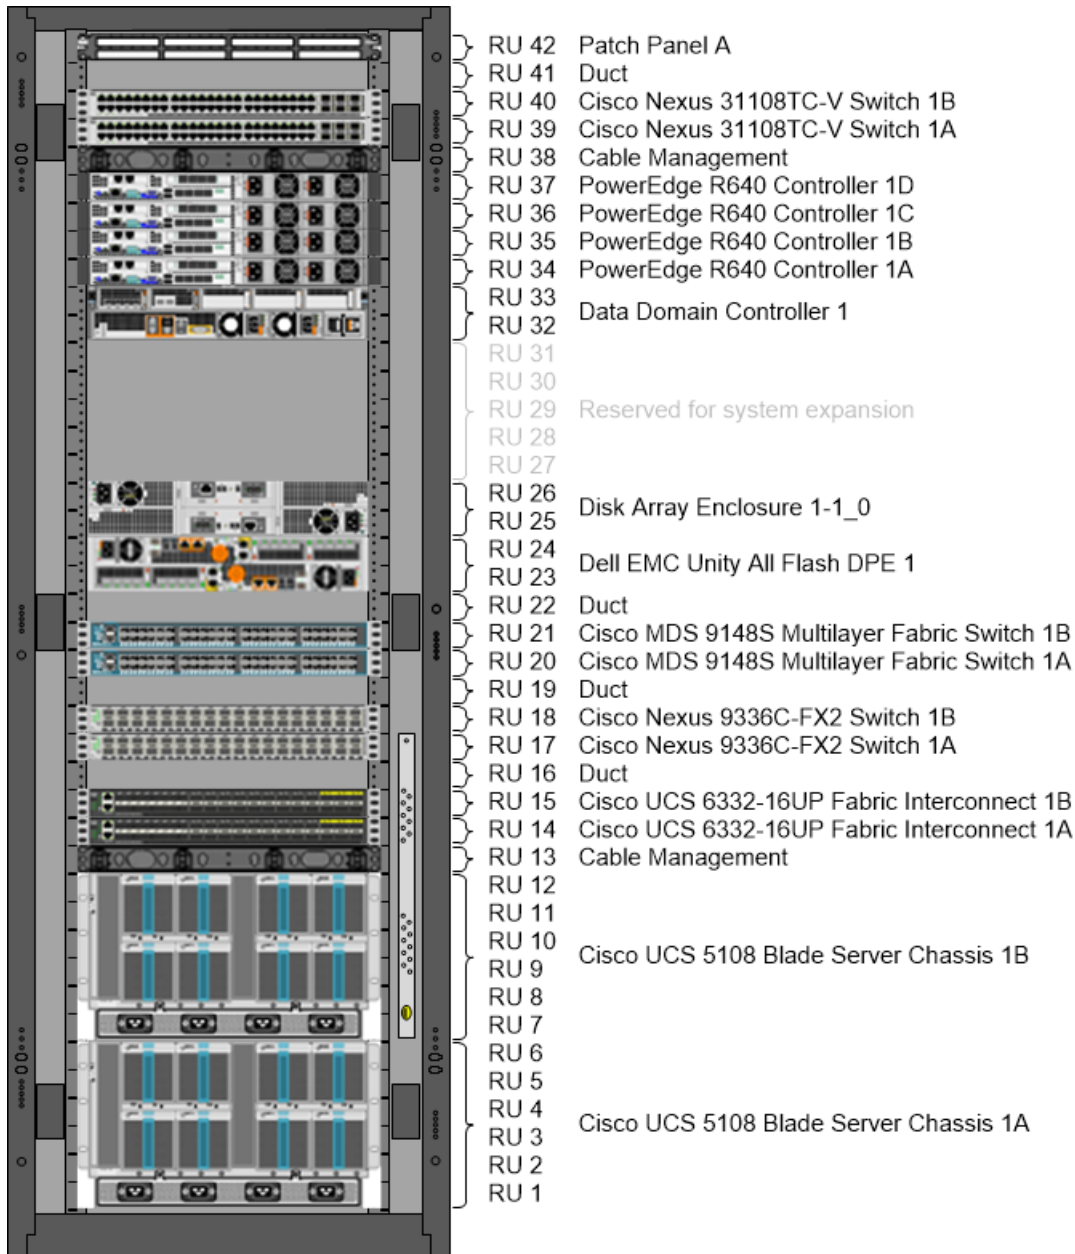


Cabinet 4



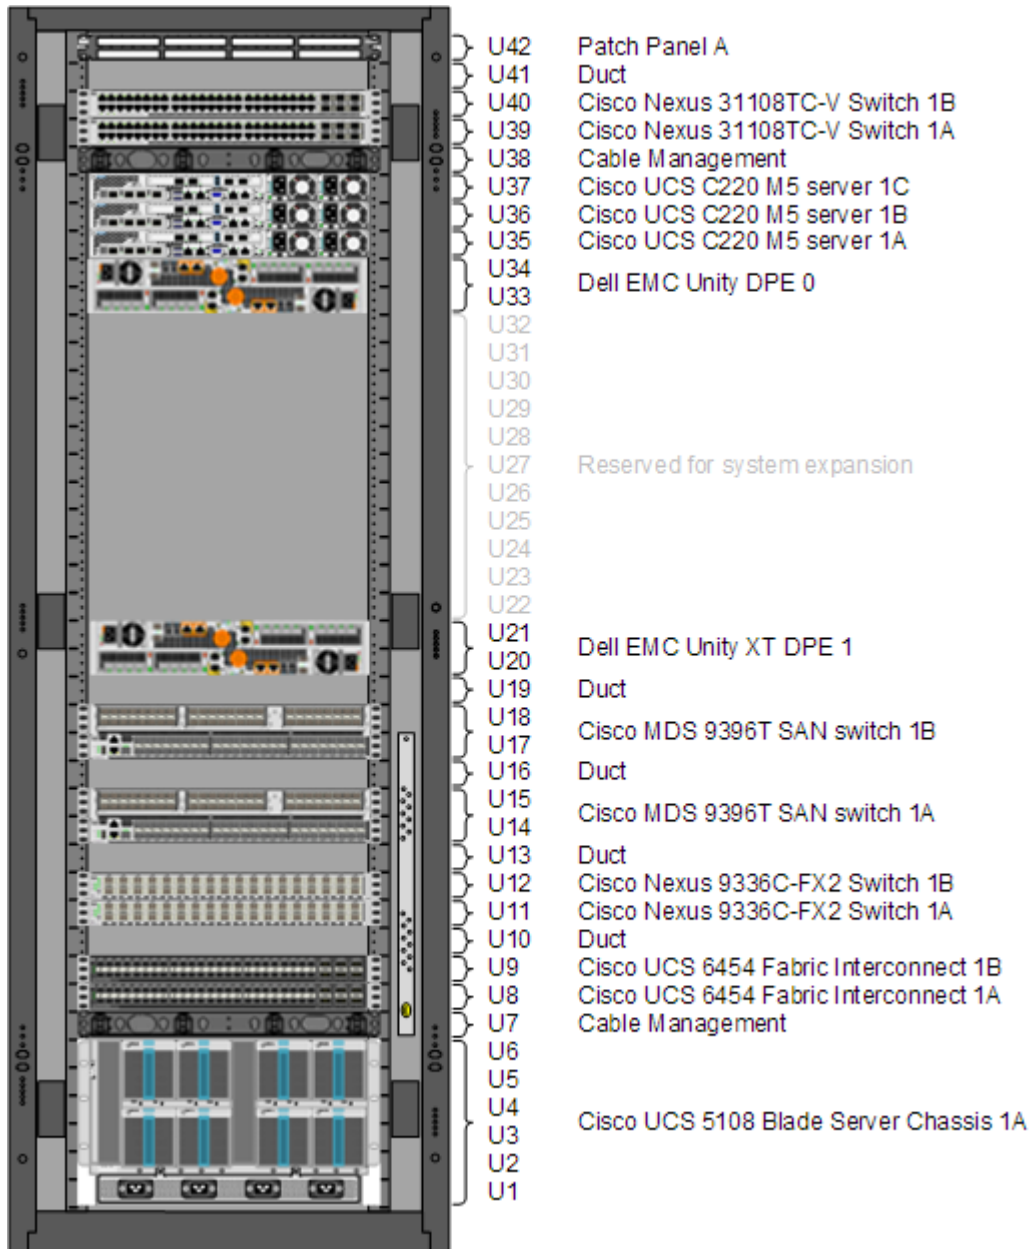
VxBlock System 1000 with Dell EMC Unity 350F storage

Elevations are provided for sample purposes only. For specifications for a specific design, consult your Dell EMC Sales Engineer.



VxBlock System 1000 with Dell EMC Unity XT storage

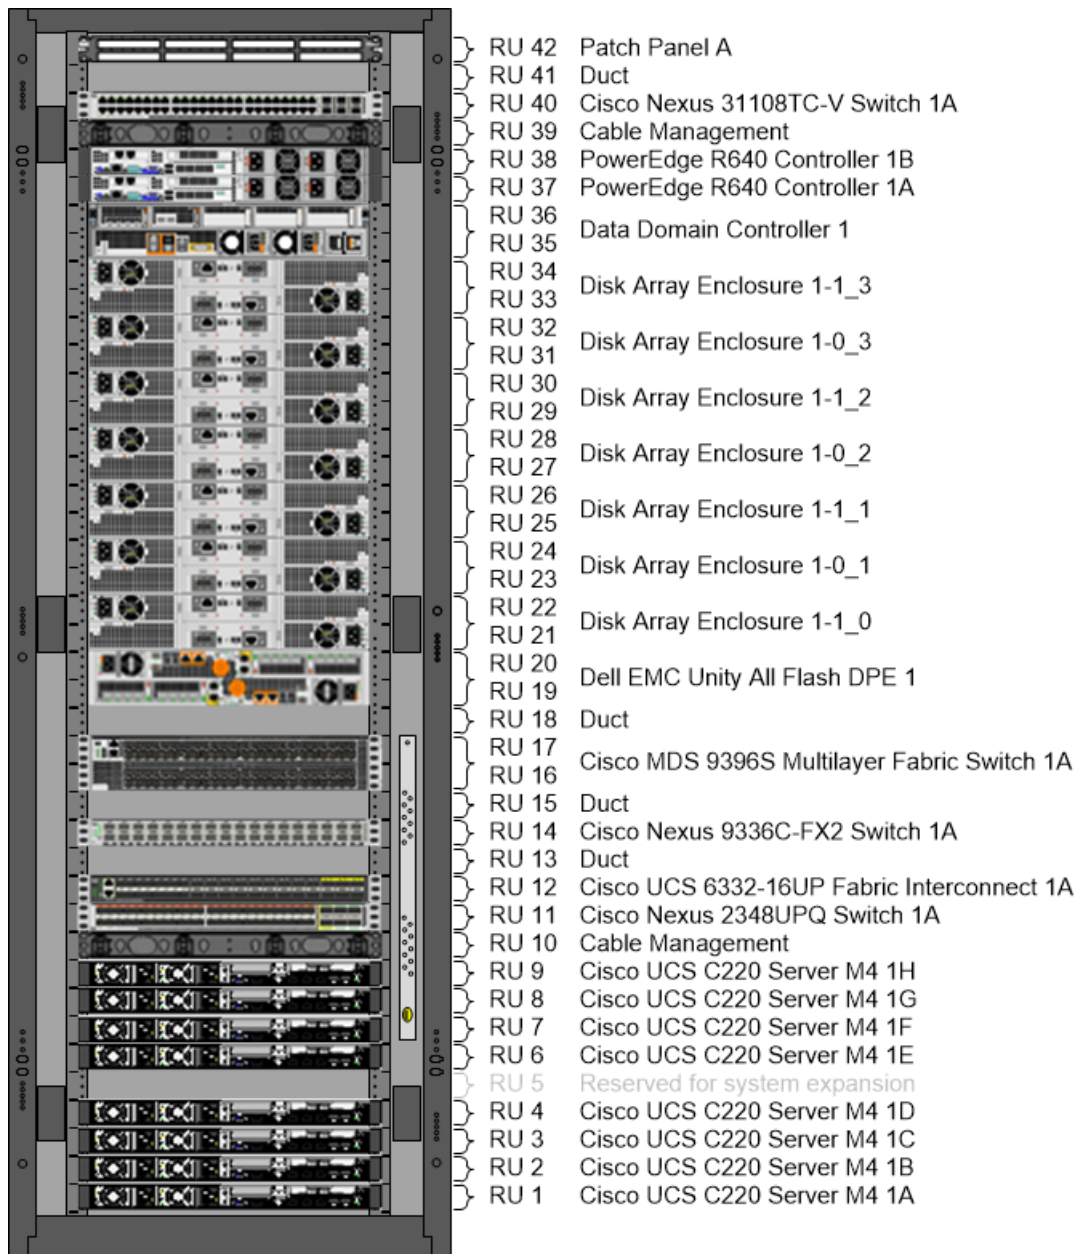
Elevations are provided for sample purposes only. For specifications for a specific design, consult your Dell EMC Sales Engineer.




VxBlock System 1000 with Dell EMC Unity and Isilon Gen 6 storage

Elevations are provided for sample purposes only. For specifications for a specific design, consult your Dell EMC Sales Engineer.


Cabinet 1



Cabinet 2



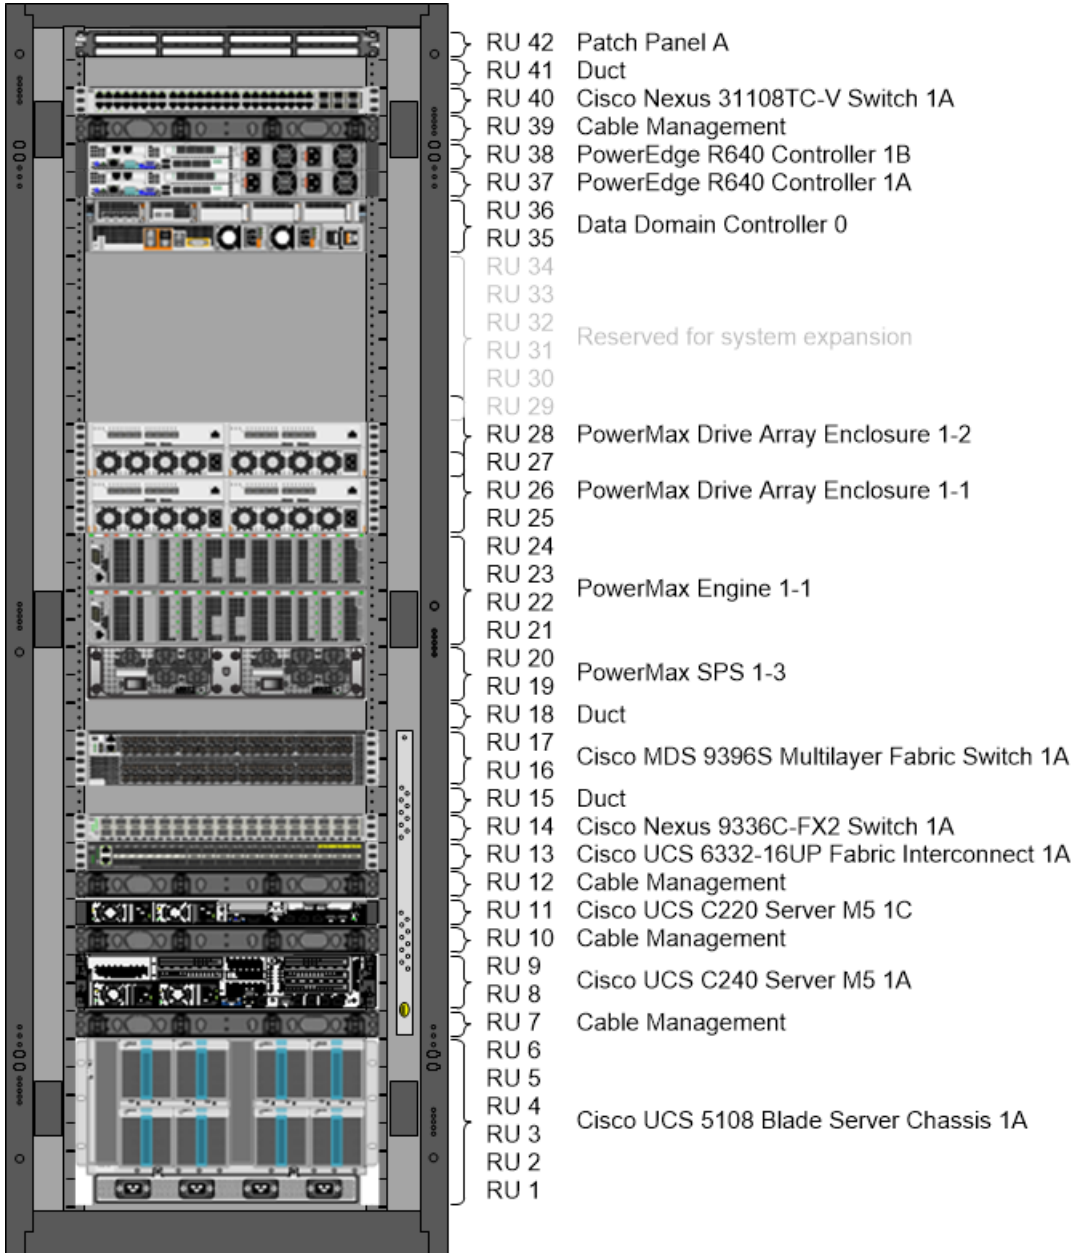
Cabinet 3




VxBlock System 1000 with PowerMax 2000 storage

Elevations are provided for sample purposes only. For specifications for a specific design, consult your Dell EMC Sales Engineer.

Cabinet 1




Cabinet 2




VxBlock System 1000 with PowerMax 8000 storage

Elevations are provided for sample purposes only. For specifications for a specific design, consult your Dell EMC Sales Engineer.

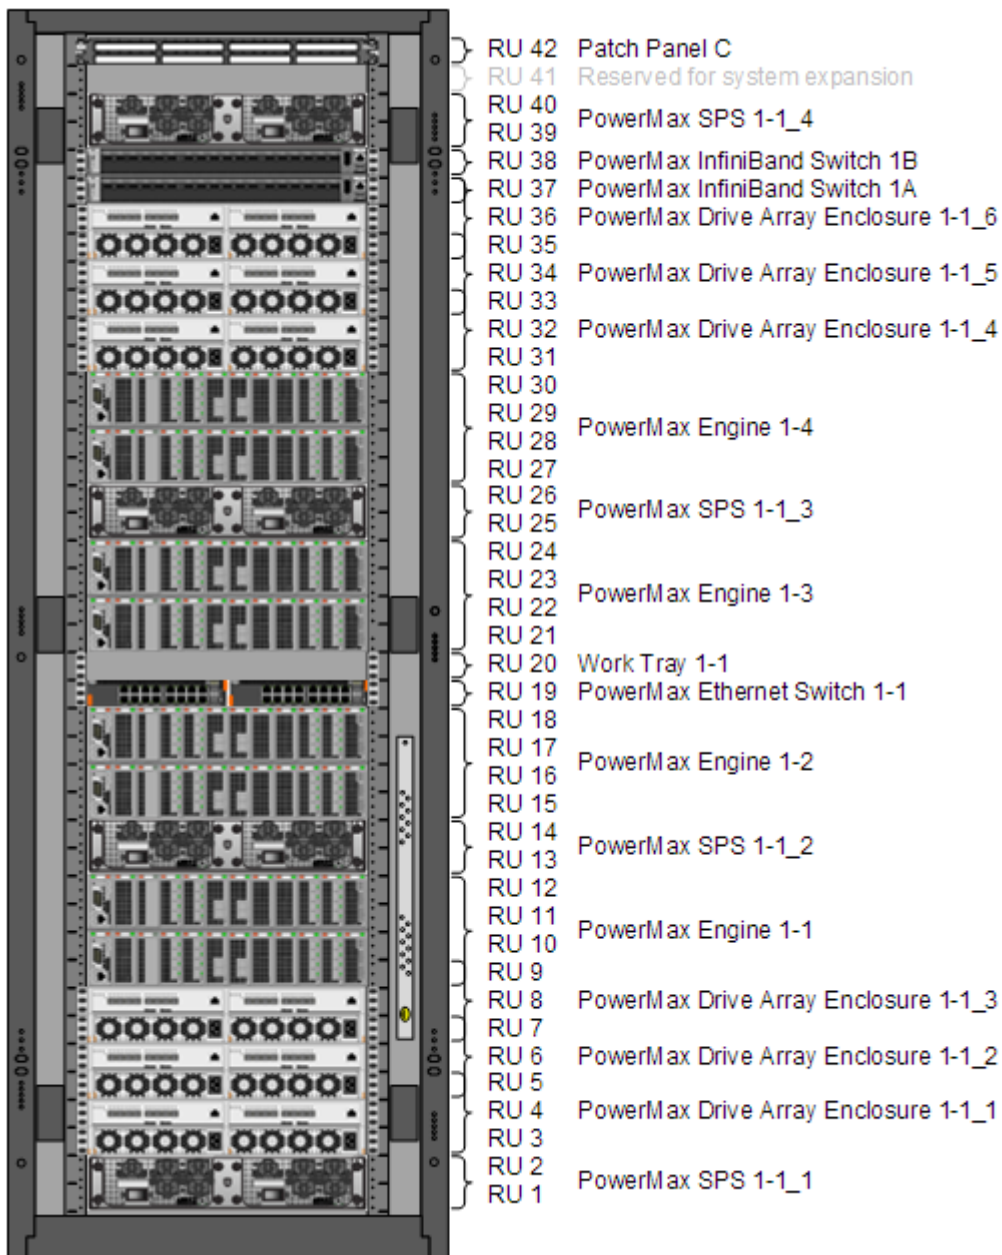
Cabinet 1



Cabinet 2



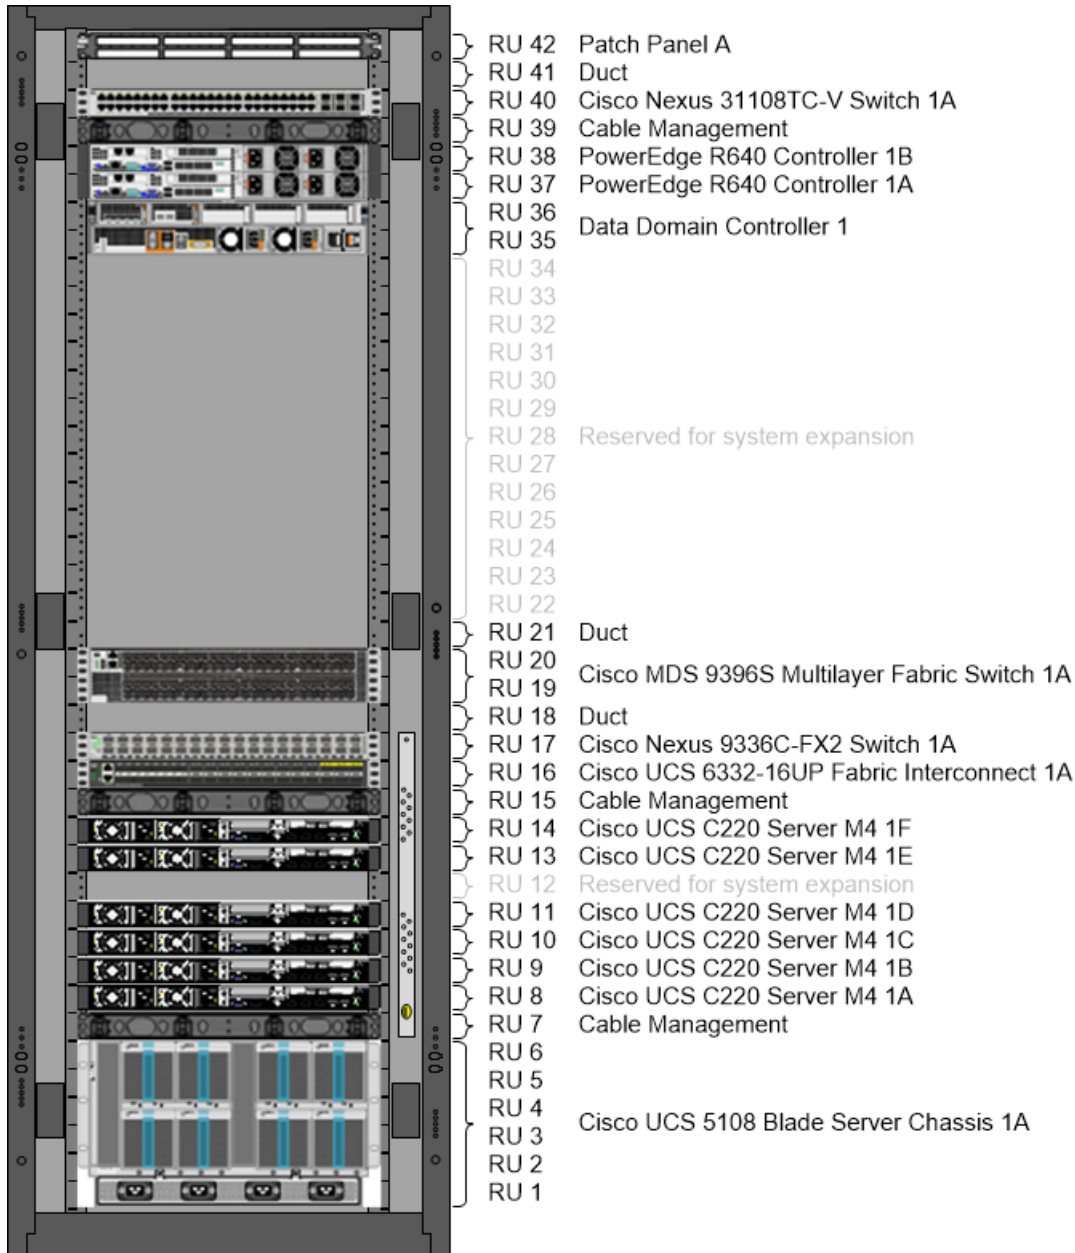
Cabinet 3



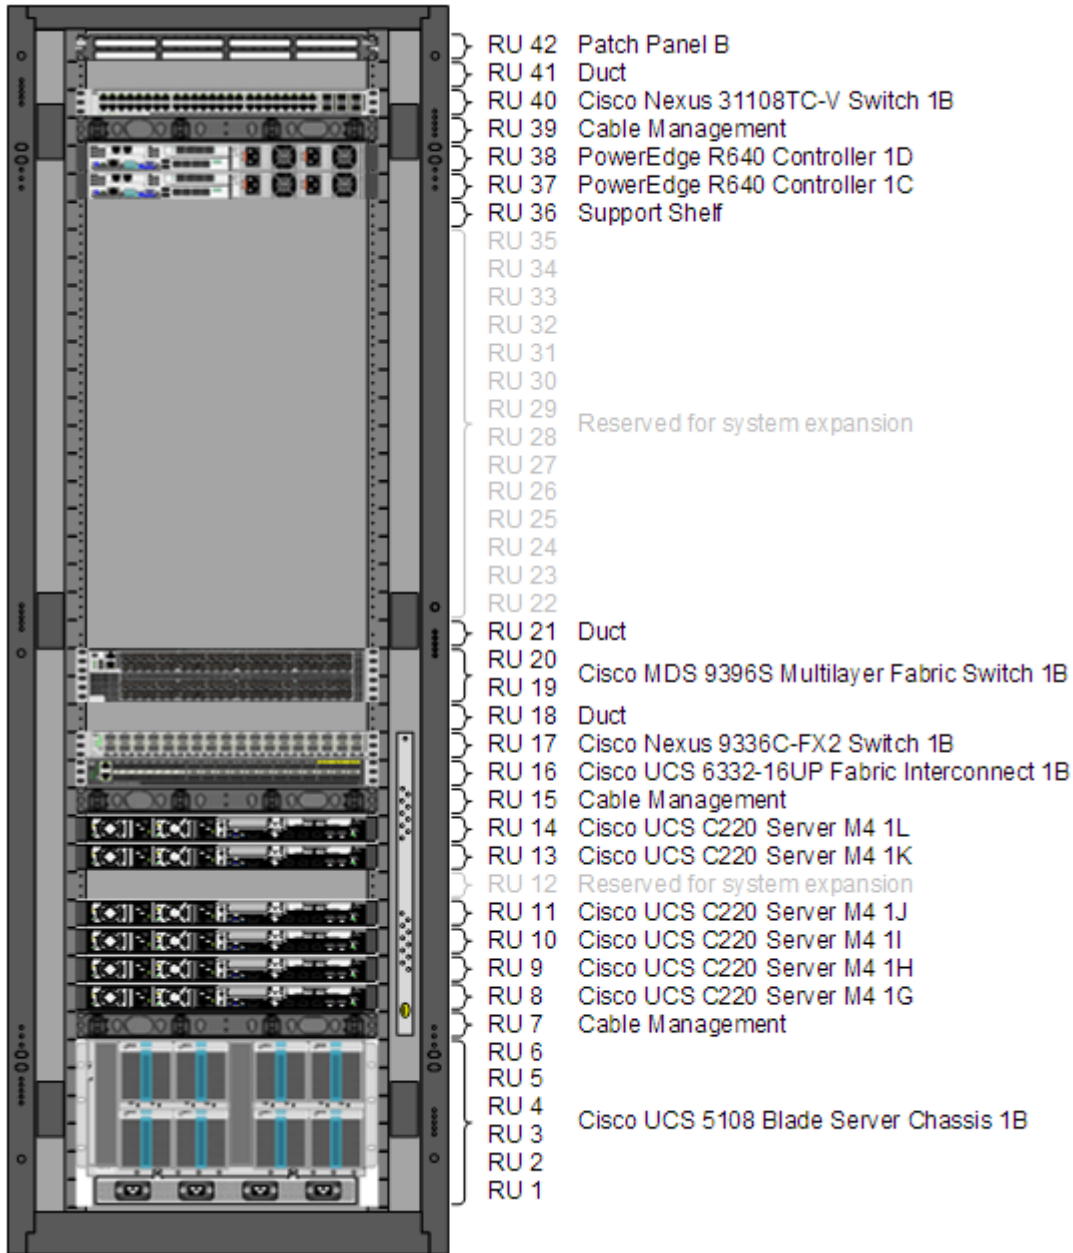
VxBlock System 1000 with VMAX 950F storage

Elevations are provided for sample purposes only. For specifications for a specific design, consult your Dell EMC Sales Engineer.


Cabinet 1



Cabinet 2




Cabinet 3




VxBlock System 1000 with XtremIO X2 storage

Elevations are provided for sample purposes only. For specifications for a specific design, consult your Dell EMC Sales Engineer.

Cabinet 1



Cabinet 2



Cabinet 3

