

Dell EMC PowerOne System

Technical Overview

Document revision 1.0

November 2019

Copyright © 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	5
Overview	6
PowerOne operational value	7
Asymmetrical scaling in PowerOne	7
Autonomous computing with PowerOne.....	8
PowerOne overview	9
PowerOne integrated architecture.....	9
Component states.....	10
PowerOne Controller.....	10
PowerOne Cluster Resource Groups.....	15
vSphere Management Cluster.....	15
Bare-metal resources.....	16
PowerOne baseline configuration and deployment.....	16
Factory assembly and baseline configuration.....	16
PowerOne deployment.....	16
Minimum and maximum configurations	17
PowerOne Controller overview.....	18
PowerOne Navigator and PowerOne API.....	18
PowerOne software upgrades.....	18
PowerOne sample use cases	19
PowerOne components	21
PowerOne compute resources.....	21
PowerEdge MX7000 chassis.....	21
PowerEdge MX740c compute sled.....	23
PowerEdge MX840c compute sled.....	23
Compute minimums and maximums.....	23
PowerOne networking resources.....	25
Dell EMC PowerEdge MX7000 networking.....	25
PowerOne storage resources.....	28
PowerMax.....	28
PowerOne network	32
System fabric.....	32
Fabric connectivity.....	32
Pod connectivity.....	33
Domain connectivity.....	33
Dell EMC Networking S5232F-ON switch to Dell EMC MX9116n Fabric Switching Engine	34
Dell EMC MX9116n to Dell EMC MX7116n.....	35
Two MX9116n configuration domain connectivity.....	35
MX7116n domain connectivity.....	36
PowerOne SAN fabric.....	38
Configuring SAN traffic.....	39
PowerOne management network.....	39
Physical architecture.....	39
Configuration of the management network.....	40
PowerOne platform organization.....	40

VMware Validated Design (VVD) considerations.....	40
PowerOne storage.....	42
PowerMax.....	42
PowerMax Integration.....	43
PowerMax Management.....	44
PowerMax storage provisioning standards.....	45
Scalability.....	45
Security.....	46
Data protection.....	46
Data Domain connectivity.....	47
Additional references.....	49
Appendix A - PowerEdge MX740c compute sled.....	50
PowerEdge MX740c processor options.....	50
PowerEdge MX740c memory.....	51
PowerEdge MX740c storage.....	53
PowerEdge MX740c PCIe networking adapters.....	53
PowerEdge MX740c supported operating systems.....	53
Appendix B - PowerEdge MX840c compute sled.....	54
PowerEdge MX840c processor.....	54
PowerEdge MX840c memory.....	55
PowerEdge MX840c storage.....	56
PowerEdge MX840c PCIe networking adapters.....	57
PowerEdge MX840c supported operating systems.....	57

Revision history

Date	Document revision	Description of changes
November 2019	1.0	Initial release

Overview

This document provides a technical overview of PowerOne. It introduces the hardware, software, and concepts on which the overall system is built. The target audience for this document includes customers, sales engineers, field consultants, and advanced services specialists.

Traditional approaches to building, operating, and maintaining data center technology, with fragmented technology silos and vendors, are no longer effective in meeting the growing needs of a modern digitally fueled enterprise. A new approach to IT and data center infrastructure is required: one founded on modern principles of operational simplicity, integration, and automation.

The Dell EMC PowerOne System is specifically designed to meet the demands of the modern, digitally driven organization. PowerOne represents the autonomous infrastructure system and Dell Technologies cloud solution built exclusively with industry-leading Dell EMC products. Although PowerOne consists of independently available components (server, storage, networking, and software), it is engineered, manufactured, managed, and supported as a single product. The Dell EMC factory team expertly assemble, integrate, and tune your system based on industry best practices before delivery, and they refine it to your exact specifications on-site.

PowerOne architecture enables asymmetrical scaling that allows you to invest only in the required resources to scale compute or storage capacity and performance independently. For example, you can reduce or increase compute and storage resources independently of each other, according to your needs. This can result in substantial cost avoidance (in hardware and software licensing) compared to more rigid (symmetrically scaling) architectural alternatives.

The highly available PowerOne Controller is the core of the system, designed to remove operational complexity. The PowerOne Controller provides a *single system* administrative experience through its centralized intelligent automation capability. This technology has been designed exclusively for PowerOne.

Unlike traditional approaches to system management and automation, the PowerOne Controller is not a tool to build, maintain, and operate your own custom-built automation. Instead, it comes *out of the box* with pre-embedded automation, enabling you to orchestrate most of the components of PowerOne based on administrative direction. This paradigm shift from *tool* to *declarative controller* enhances operational efficiency, simplicity, and agility. The PowerOne Controller helps you to avoid the manual burden of researching, qualifying, and running many procedures and integrations that often represent the largest cost of running a data center.

As the orchestration layer for PowerOne, the PowerOne Controller exposes its full capability through an Application Programming Interface (API) for simple business process integration and on-premise cloud-like operations. In addition, PowerOne includes a modern and intuitive user interface (PowerOne Navigator) for PowerOne operations.

PowerOne operational value

Based on your user-defined goals, PowerOne delivers intelligent automation that matches these goals to available resources and configures the system to best meet those goals.

The benefits of PowerOne are:

- Reduces Total Cost of Ownership (TCO) by automating routine operations associated with enterprise-class infrastructure.
- Accelerates time to value for production workloads by reducing the time from purchasing the system to production readiness.
- Reduces operational risk by focusing human inputs towards goals and automating the configuration that meets the requirements of those goals.
- Reduces stranded components, by driving up their utilization, providing a better Return On Investment (ROI), and eliminating the risk associated with moving underused assets to other services.
- Provides continuous configuration validation and lifecycle management for all hardware and software components.
- Reduces operational risk and downtime with monitoring and alerting, including call home for all components.

PowerOne provides the ease-of-use of a public cloud experience, while retaining the control, safety, and optimization of an on-premises system, by:

- Avoiding latency issues.
- Detecting potential cost overruns and underuse of resources.

Asymmetrical scaling in PowerOne

In PowerOne, disaggregated infrastructure resources, such as compute and storage, scale up and scale out independently. This asymmetrical scaling means that you do not have to grow the system in fixed proportions, driving efficiency and better resource utilization. PowerOne enables you to define your resources when you order, and expand and fine-tune them on an ongoing basis.

With PowerOne you can:

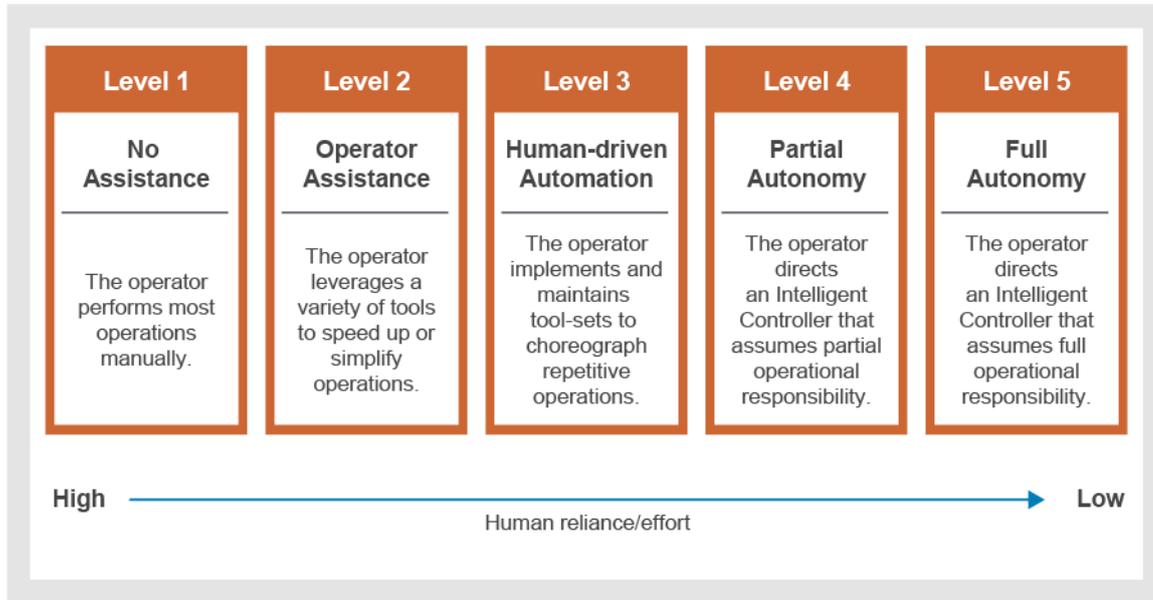
- Scale up by adding more capacity
- Scale out by adding more arrays and or servers and or extending the network fabric
- Scale storage capacity independent of compute capacity
- Scale storage performance independent of storage capacity

The PowerOne Controller software, created specifically to support asymmetrical scaling, builds upon and extends this underlying flexibility. You can create cluster resource groups (CRGs), each with a specific resource profile that has the right mix of resources to support your applications' performance and capacity needs. CRGs can scale asymmetrically by configuring the right amount of resources for a particular application as required. An ecosystem of resource groups with optimum proportions can then be scaled out to deploy additional applications and application types as business needs expand.

For more information on CRGs, see [PowerOne Cluster Resource Groups](#).

Autonomous computing with PowerOne

Based on a wider view of data center infrastructure, it is observed that true autonomous computing is largely absent from data centers today. You can gauge the level of automation in your data center by comparing it to the following taxonomy:



PowerOne provides Level 4 capability by offering tools and methods to automate processes that previously required manual intervention.

Reliable automation enables you to perform what you do now on a larger scale, at greater efficiency, and with less downtime—by automating routine tasks that previously required a high degree of human intervention.

PowerOne overview

PowerOne is an integrated infrastructure solution that provides autonomous provisioning of system resources. The system is designed as a single unit that can be easily expanded after initial deployment.

PowerOne comprises the following system resources:

Compute

PowerOne incorporates the Dell EMC PowerEdge MX compute platform to deploy compute resources as virtualized or bare metal.

Virtualization

PowerOne delivers an optimized integrated offering for VMware configurations. It also enables your application deployments on bare-metal servers with native operating system installations.

Fabric

PowerOne uses infrastructure fabric that allows for seamless scalability from small to large deployments. Dell EMC networking resources ensure that PowerOne offers a modern network fabric that is designed to simplify and automate PowerOne fabric implementation and management.

Storage

PowerOne uses Dell EMC PowerMax storage arrays with Non-Volatile Memory Express (NVMe) storage technology, providing an all-flash storage platform designed for demanding environments.

Data Protection

PowerOne allows integrated data protection resources, such as Avamar/Data Domain, NetWorker/Data Domain, and RecoverPoint for Virtual Machines (RP4VM).

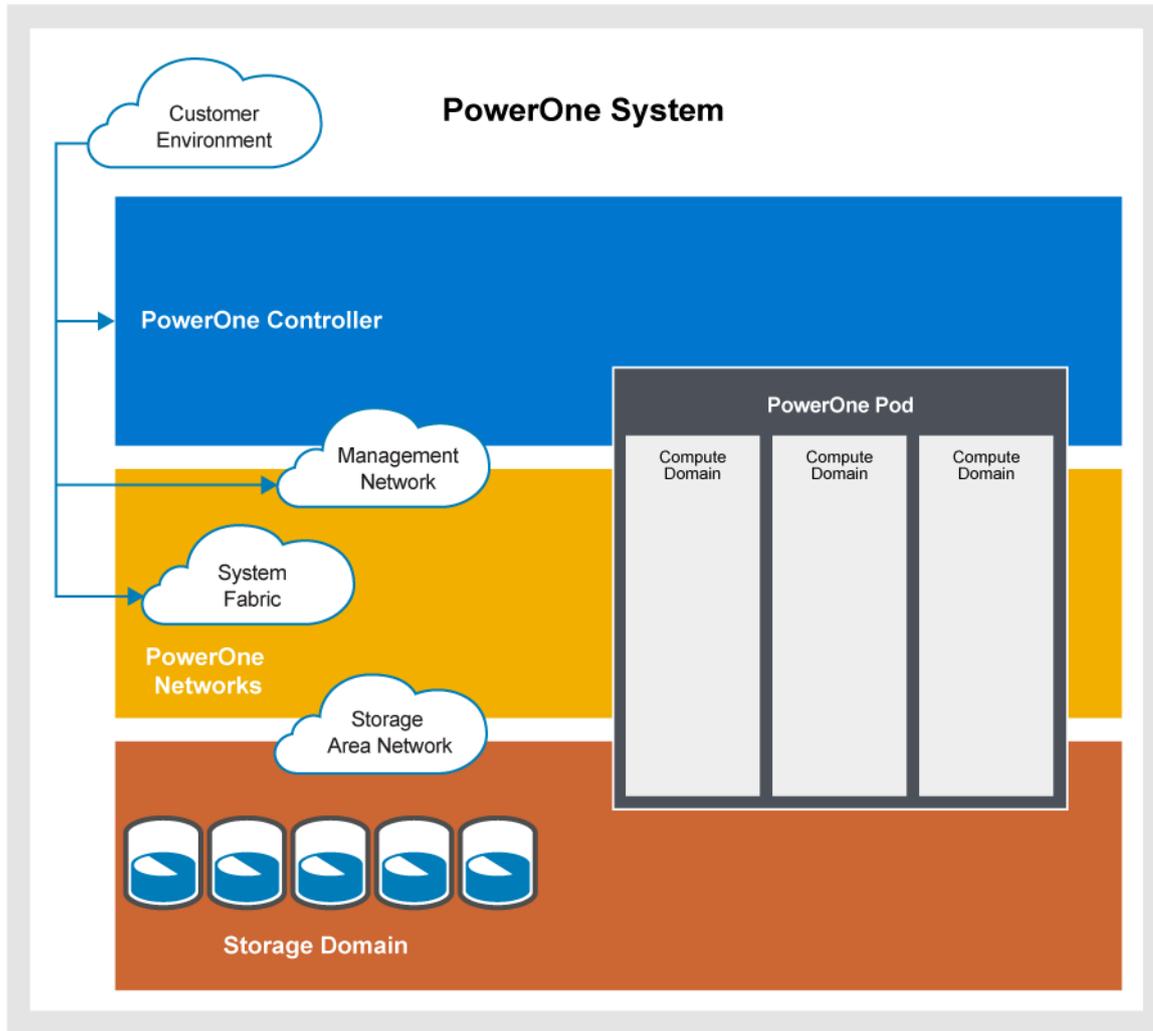
Management

PowerOne includes the PowerOne Controller, a management resource that directs the PowerOne components and regulates them through the PowerOne Navigator user interface (UI) and the PowerOne Controller API.

PowerOne integrated architecture

The following PowerOne components are assembled to provide an integrated system architecture:

- PowerOne Controller
- Compute domains
- PowerOne Pod
- Storage domain
- PowerOne networks



Component states

Each component in PowerOne has a component state that describes its operational status. The component state determines the actions that are available for the components.

After a workflow that configures a particular component is completed in PowerOne, the state of that component is updated automatically. PowerOne also updates component states during regular inventory scans.

In PowerOne Navigator, the system detects any component in the `New Asset` state, and allows you to apply a baseline configuration to the component that changes its component state to `Automatable`.

In PowerOne Navigator you can change the compute sled's firmware if the component is **Automatable** and if that compute sled is not part of a cluster.

PowerOne Controller

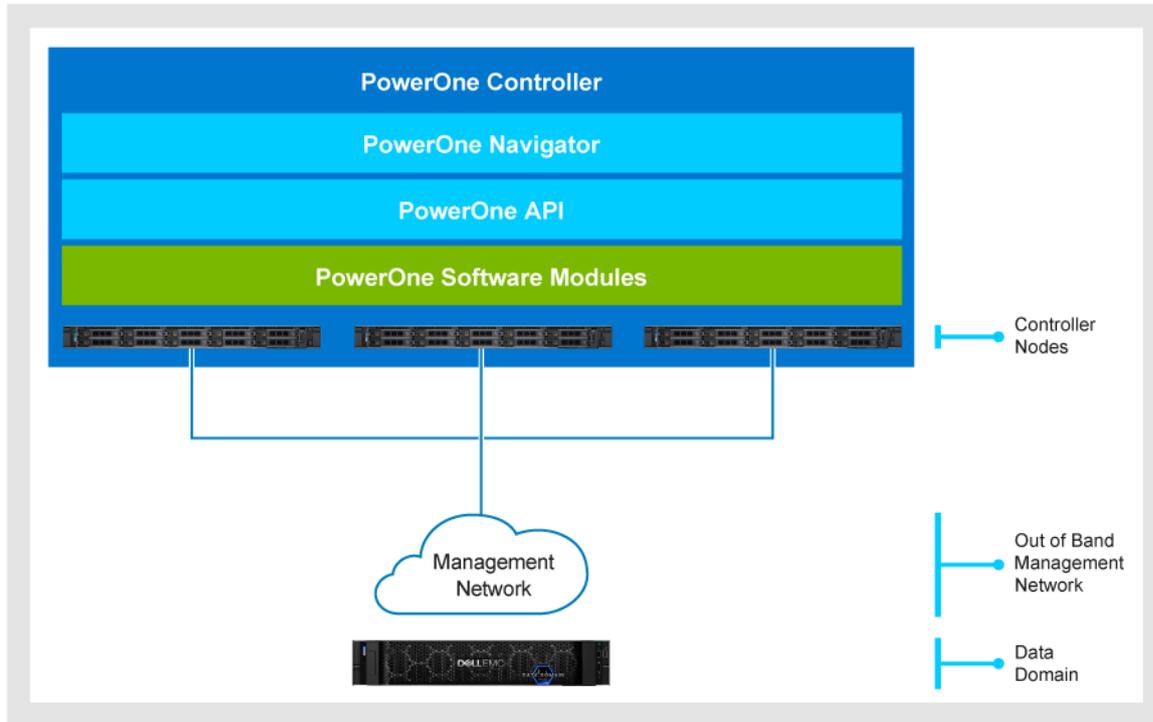
The PowerOne Controller performs several operations that manage PowerOne. It provides a dedicated unified management plane for PowerOne compute domains, PowerOne Pods, PowerOne networks, and PowerOne storage domains. It orchestrates these components together into a unified group of PowerOne components known as a Cluster Resource Group (CRG).

The PowerOne Controller is a dedicated system reserved solely for preinstalled PowerOne management workloads. It is a highly-available platform that hosts the PowerOne Controller software. It consists of three PowerOne Controller rack

servers that are clustered to create a single unified platform for using the PowerOne Controller software. Each of the three PowerOne Controller servers contains flash storage.

The PowerOne Controller nodes are directly accessed only by the PowerOne API or PowerOne Navigator.

The following figure shows the PowerOne Controller and the PowerOne management network infrastructure.



Data Protection

Dell EMC NetWorker provides data protection for the PowerOne Controller. The data protection system uses two separate virtual machines (VMs) to protect the PowerOne Controller. One VM is deployed in the PowerOne Controller (called Bastion), the other VM is a NetWorker VM (called NetWorker Server) deployed on the vSphere Management Cluster.

Bastion is responsible for collecting backups of the PowerOne Controller and transmitting them to the NetWorker Server. The NetWorker Server has a 1 TB VMDK virtual disk (Eager Thick provisioned) on a VMFS volume in the vSphere Management Cluster. The VMFS volume is provisioned on the PowerOne storage domain. This configuration creates separate production and backup volumes, which allows the PowerOne controller to recover in case of failure.

NetWorker performs full-image level backups of the PowerOne Controller daily.

Backups of the PowerOne Controller are retained for three days, assuming you perform a single backup each day. If you require longer retention periods or more frequent backups, you can increase the NetWorker license from 1 TB to the appropriate size for your data retention needs.

Optionally, you can purchase Data Domain to expand your data protection capacity for the PowerOne Controller. If you select Data Domain to protect production workloads in PowerOne, you can use the same Data Domain to protect the PowerOne Controller. This reduces the number of data protection devices in the environment.

PowerOne Pod

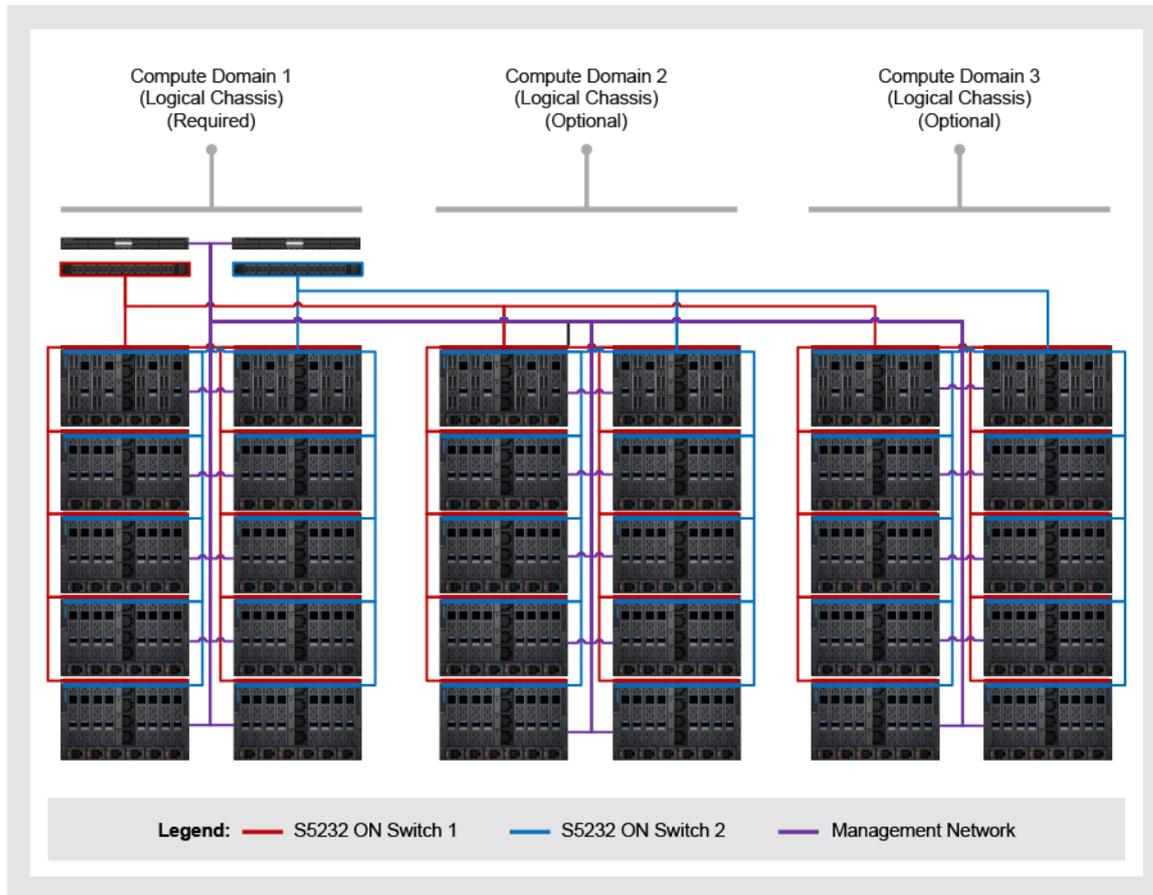
The PowerOne Pod is a collection of components, physically separated into compute domains that connect northbound to a single pair of S5232F-ON switches.

A PowerOne Pod consists of the following:

- A maximum of three compute domains. Each compute domain contains a minimum of eight and a maximum of 80 compute sleds. A maximum of 240 compute sleds are available per PowerOne Pod.

- A maximum of 30 MX7000 chassis are divided into three compute domains.

The following figure shows a PowerOne Pod consisting of three compute domains, each having two S5232-ON edge switches and two OOB management S4148T-ON switches. It also shows the interconnections between the MX7000 chassis.



Related information

[PowerOne networking resources on page 25](#)

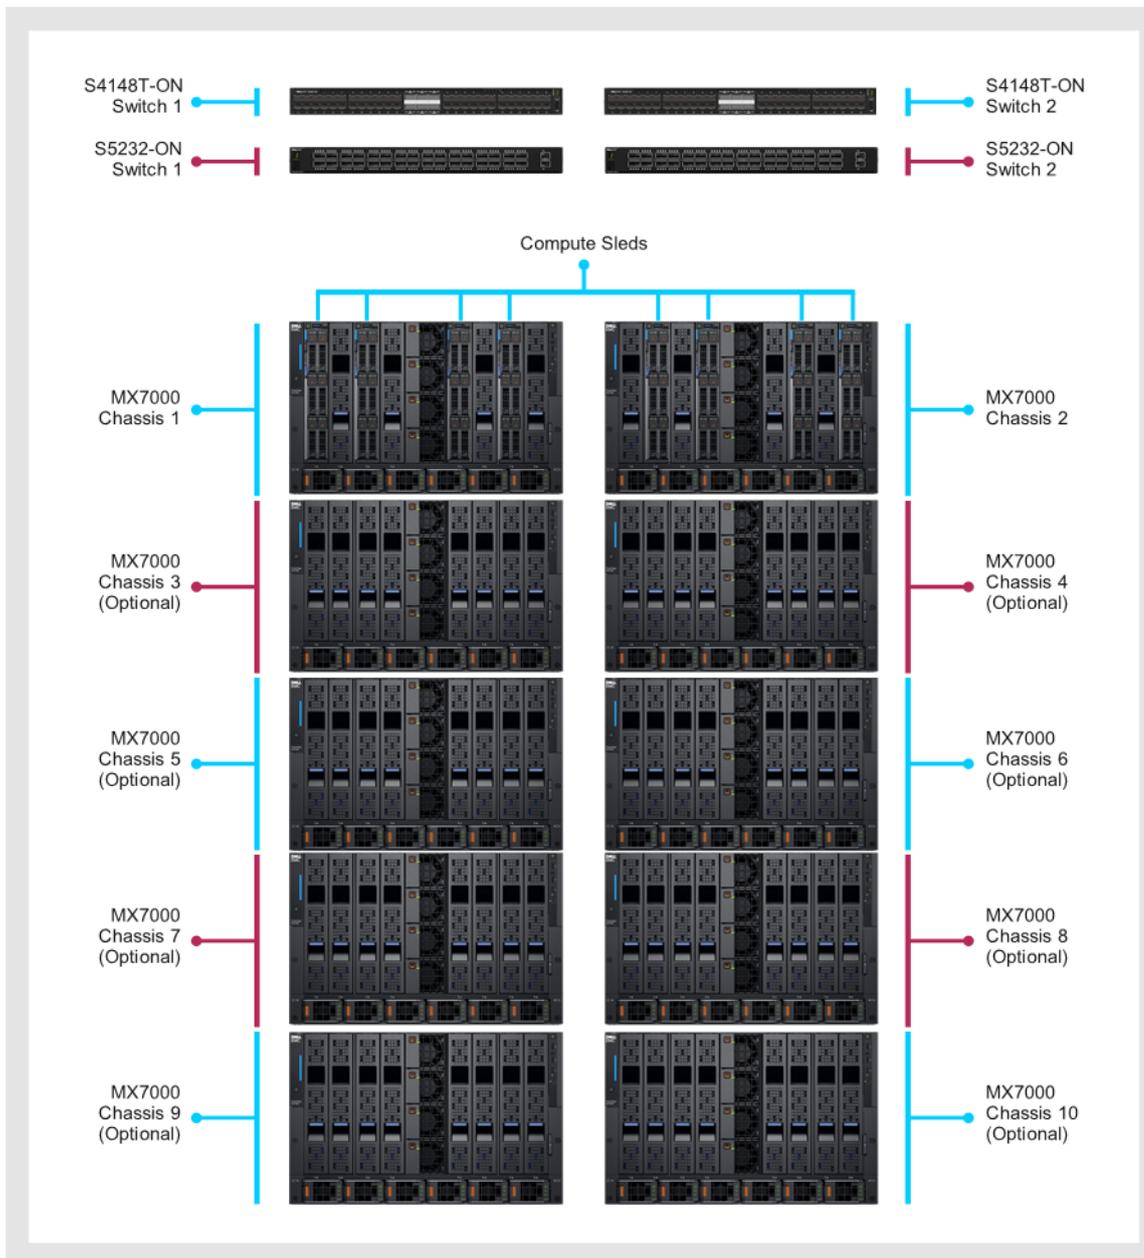
Compute domain

One of the core components of a CRG is its computing capacity. Computing capacity is distributed among logical blocks known as PowerOne compute domains. A compute domain is a collection of interconnected Dell EMC PowerEdge MX7000 chassis in PowerOne. Each compute domain consists of a minimum of two and a maximum of ten MX7000 chassis. The first two chassis in a compute domain are designated as primary chassis and have one pair of Dell EMC MX9116n Fabric Switching Engines that provide PowerOne System Fabric connectivity. Each of the remaining eight MX7000 chassis connects to the primary MX7000 chassis using the Dell EMC MX7116n Fabric Expander Modules.

Each PowerOne Pod contains a pair of OOB management S4148T-ON switches and two S5232-ON Leaf switches. These OOB and Leaf switches form the boundaries of a PowerOne Pod. The example compute domain in the following figure shows:

- A pair of S4148T-ON out-of-band (OOB) management switches
- A single compute domain with two S5232-ON Leaf switches
- Two MX7000 partially populated with four MX740c compute sleds in each chassis

- Eight optional MX7000 with sled blanks



The compute domain in the example provides a total capacity of 80 two-CPU (MX740c) compute sleds.

A minimum of one compute domain is required within PowerOne, and each PowerOne Pod must contain at least one compute domain. PowerOne Pod is a collection of one to three compute domains, and associated management and network switches.

Compute domains scale at the MX7000 chassis level, starting at a minimum of two MX7000 chassis to a maximum of ten chassis per compute domain. The MX7000 chassis scale at the compute sled level and can hold up to eight MX740c compute sleds or four MX840c compute sleds.

Only two MX7000 chassis are required for each compute domain. The first compute domain in PowerOne must be populated with a minimum of four compute sleds in each of the first two chassis, for a total of eight compute sleds across the two chassis. Four compute sleds are used for the CRG that hosts the vSphere Management Cluster, which contains management applications such as VMware vCenter. The second CRG created consists of four or more compute sleds and represents the workload cluster in a VMware Validated Design (VVD).

Storage domain

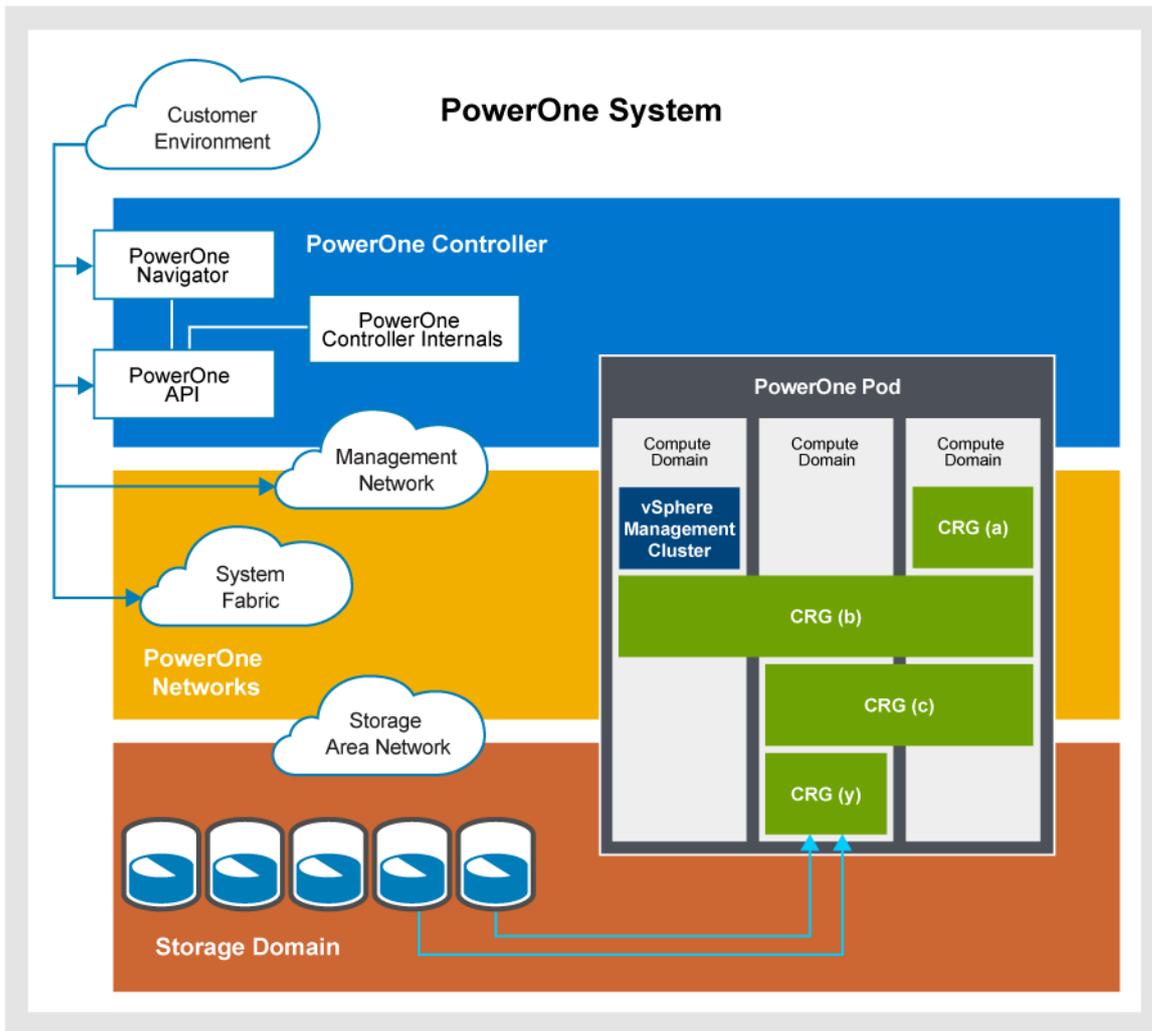
The storage domain provides production workload storage from managed arrays to one or more compute domains. The storage domain is a collection of supported storage components of PowerOne that are automated by the PowerOne Controller.

The storage domain supports an all-flash storage and consists of a minimum of one PowerMax array.

When you create a CRG, PowerOne allocates the appropriate storage (capacity and SAN paths) from the available storage resources. PowerOne then maps that storage volume to the compute hosts in the CRG. This process is repeated when additional storage is added to the CRG.

The following figure shows two volumes from the storage domain that are mapped to CRG (y). This is important because the storage and compute resources are being mapped together at the physical level. This mapping occurs before the CRG resources are managed by a hypervisor or an operating system. A unique function of PowerOne is to map the components to a CRG before the deployment of a hypervisor.

This grouping of components at the infrastructure level enables consumption of disparate resources as a single logical unit, a CRG. Thus, even before a hypervisor is applied, a logical mapping is performed on the various components consumed. For example, as a CRG is expanded by adding more compute capacity, it is not necessary to manually map existing volumes to the newly added compute resources; PowerOne performs this task.



PowerOne Cluster Resource Groups

A Cluster Resource Group (CRG) is a flexible collection of compute, network, storage, and software resources.

To provide predictable performance, even during hardware failures, all servers in a CRG share the same firmware, configuration, and BIOS settings.

PowerOne includes the types of CRGs listed in the following table.

CRG Type	Description
VMware Workload Cluster	The workload cluster is a CRG instantiated on at least four compute sleds in the first compute domain following the instantiation of the vSphere Management Cluster. This CRG hosts tenant VMs (sometimes referred to as workloads or payloads). This shared cluster also runs the VMware NSX services to enable north-south routing between tenant virtual machines and the external network, and east-west routing inside the PowerOne. As the environment expands, additional CRGs can be created to support a mix of different types of workloads and Service Level Agreements (SLAs).
vSphere Management Cluster	A vSphere Management Cluster is instantiated by the PowerOne Controller on a vSphere Management Cluster consisting of four compute sleds in the first compute domain. This cluster lives in the management workload domain and runs the management virtual machines. These virtual machines include vCenter Server, vSphere Update Manager, NSX Manager, NSX Controller, vRealize Operations Manager, vROps Management Pack, and other management components both deployed by PowerOne and by customers. These VMs are protected by VMware High Availability (HA). For more information, see the vSphere Management Cluster section.

A CRG can span multiple compute domains within the limits of the controlling platform. For example, in a VMware vSphere 6.7 vCenter Server environment, if the maximum cluster size of compute components is 64 nodes, the maximum size of a VMware-based CRG is 64 compute sleds.

Because compute sleds in CRGs can span PowerOne compute domains, placement of hosts is not architecturally constrained.

After you create a CRG, deploy the operating systems or a hypervisor on the hosts in the CRG to create a cluster of resources on which the workloads are deployed. After a CRG is provisioned, the PowerOne Controller deploys the VMware ESXi virtualization components on the hosts in the CRG.

Note: All the servers in a CRG must share the same firmware, configuration, and BIOS settings during normal operations.

CRGs differ based on the following:

- **Production workload types:** Workloads from development labs, virtual desktops (VDI), cloud, database, production operations, and so on are different.
- **Hardware life cycles:** Life cycles of each component vary based on component depreciation.

vSphere Management Cluster

The first CRG in PowerOne, the vSphere Management Cluster, consumes a minimum of, four compute sleds in the first compute domain. The vSphere Management Cluster runs management VMs that control the hypervisors (ESXi hosts) on both vSphere Management Cluster and other clusters.

Management virtual machines can include both—those deployed by the platform (such as vCenter) and those that are deployed by the operator.

Provision management, monitoring, and infrastructure services on the vSphere Management Cluster, which is configured as a VMware vSphere cluster that provides high availability to these critical services.

The vSphere Management Cluster has a minimum of two VMware vCenters: a management vCenter and a compute vCenter. The management vCenter controls the virtual environment of the vSphere Management Cluster. The compute vCenters manages the virtualization of one or more CRGs.

The vSphere Management Cluster also contains VMware NSX-v components, which can be optionally configured to manage NSX-T deployments for customer workloads (CRGs). In PowerOne, the PowerOne System Fabric provides network connectivity from the customer environment to the vSphere Management Cluster.

VMware vRealize Operations (vROps) is optional for use in PowerOne. If it is purchased as part of PowerOne, vROps management packs are included. The Dell EMC PowerOne management packs are developed for vROps. They monitor the health, capacity, and performance of system resources in PowerOne, and help to rapidly identify and troubleshoot system issues.

Bare-metal resources

PowerOne allows bare-metal resource allocation. They are independent resources that are not managed or controlled by PowerOne.

Compute sleds can be made available for bare-metal allocation by setting them to manual mode. Storage must be provisioned through port groups reserved for bare metal.

PowerOne baseline configuration and deployment

PowerOne is assembled at the Dell EMC factory and tested before shipment. On arrival at the deployment site, the customer defines a production configuration for their PowerOne.

Factory assembly and baseline configuration

The factory assembly process includes physical racking and cabling of the components. After all the components are assembled and tested, a baseline configuration is applied.

The baseline configuration contains an inventory of PowerOne components and the state of each component. The baseline configuration is generated and stored in the PowerOne Controller. After the baseline configuration is complete, PowerOne is packaged for shipping to the deployment site. On the arrival at the deployment site, powering on the PowerOne Controller begins the deployment process.

PowerOne deployment

After PowerOne is installed in the data center, powering on the PowerOne Controller begins the deployment process. Then, the PowerOne Controller verifies to ensure that all components are present, and nothing has been damaged or misplaced during shipping.

The PowerOne production configuration starts after the PowerOne Controller is powered on. The PowerOne Navigator Deployment Wizard displays software activation and license terms for your review and acceptance. The PowerOne Deployment Wizard then guides you through integration with existing customer networks, with assistance from Field Services.

After the infrastructure is configured, PowerOne deploys a VMware vSphere Management Cluster on four of the compute sleds across the first two chassis in the first compute domain. The vSphere Management Cluster houses the virtualization management resources. This completes the initial deployment of PowerOne.

After PowerOne is connected to the data center network, it is ready to provision CRGs and their associated vSphere clusters. For more information on CRGs, see the [PowerOne Cluster Resource Groups](#) section.

Minimum and maximum configurations

The following table lists the minimum and maximum configurations for a PowerOne.

Component	Minimum	Maximum	Notes
PowerOne Controller	1	1	
PowerOne Controller nodes	3	3	Not customer/partner/Sales Engineer (SE) expandable.
S5232-ON switches per Pod	2	2	
Compute sled slots per Pod	8 (4 x 2 x 1)	240 (80 x 3)	
Compute domains per Pod	1	3	
MX7000 chassis per compute domain	2	10	
Compute sled slots per compute domain	8	80	
Compute sleds in first compute domain	8	80	Combining MX840c compute sleds with MX740c compute sleds in a single compute domain may reduce the overall sled capacity of the compute domain.
MX740c compute sleds in second and third compute domains	0	80	Combining MX840c compute sleds with MX740c compute sleds in a single compute domain might reduce the overall sled capacity of the compute domain.
MX840c compute sleds in second and third compute domains	0	40	Combining MX840c compute sleds with MX740c compute sleds in a single compute domain reduces the overall sled capacity of the compute domain.
MX7000 Chassis per PowerOne	2	30	
Compute sleds per platform (vSphere Management Cluster)	4	-	
MX740c compute sleds per PowerOne	8	240	
Arrays per storage domain	1	-	
PowerMax 2000 E/P Bricks per mixed rack	-	2	
PowerMax 8000 E/P Bricks per mixed rack	1	8	

PowerOne Controller overview

The PowerOne Controller manages and controls PowerOne, and it is responsible for management of system resources that are deployed as CRGs at the physical layer. One important distinction is that the PowerOne Controller is present outside the PowerOne Pod and storage domain.

The PowerOne Controller works with the platform to:

- Perform initial production configuration of PowerOne. For example, define IP address ranges, DNS, NTP and syslog targets, and so on.
- Provide an inventory of available system resources, including firmware and deployment disposition.
- Create CRGs or add resources to CRGs. PowerOne communicates this to vCenter.
- Create the vSphere Management Cluster that hosts vCenter.

Related information

[vSphere Management Cluster on page 15](#)

PowerOne Navigator and PowerOne API

The PowerOne Controller is accessed either through the PowerOne API or PowerOne Navigator. The PowerOne API is a fully functional REST API designed to enable users to integrate PowerOne into their existing management systems.

PowerOne Navigator is a dynamic user interface that is built on top of the PowerOne API to provide visibility into, and interact with PowerOne. PowerOne Navigator provides an intuitive GUI to initially deploy PowerOne and manage its automated day-to-day operations, such as operator accounts, CRGs, and version controls.

PowerOne software upgrades

PowerOne uses a dynamic update approach based on the VMware Validated Design (VVD) releases. As new VVD versions are released, those updates along with associated firmware, drivers, and software components are tested against PowerOne. The resulting tested hypervisor, firmware, drivers, and software are assembled in a bundle that is tied to the VVD release. The bundles are made available to update both vSphere Management Cluster and CRGs. The bundles are applied intelligently when you choose to update a vSphere Management Cluster or CRG, or create a CRG in PowerOne Navigator. Updates to CRGs are dependent on the bundle that is applied to the vSphere Management cluster.

PowerOne sample use cases

PowerOne is ideally suited to host a wide variety of business workloads. Some sample uses are described in this chapter.

Data center and workload consolidation and modernization

One common and effective IT strategy is to consolidate different workloads, with common or diverse performance requirements, in the same infrastructure assets. This consolidation strategy reduces infrastructure costs, saves space and power, and optimizes management efforts.

For this consolidation strategy to be effective, the infrastructure needs to provide ample flexibility in terms of scalability and performance, as well as provide attractive cost profiles for every workload. PowerOne is an example of a highly scalable and automated system, able to host workloads with a wide range of performance, availability, and scalability requirements.

With the highly scalable family of PowerEdge MX7000 servers and PowerMax storage systems, PowerOne can incorporate the infrastructure elements needed to host almost any type of workload, including high performance databases, dense virtualized environments, modern and mission-critical applications—all in the same platform.

In PowerOne, the rich set of data, security, and availability services are a key element necessary for data center consolidation strategies, because different workloads will have different needs in terms of data services, security, and availability of requirements.

Build a private cloud

PowerOne is the perfect platform choice for building a private cloud because it provides:

- Flexible infrastructure options that can match any virtualized workload needs in the compute, storage, and network domains.
- Design and engineering as a VMware vSphere platform with cloud management and operations integration.
- Highly automated operations that conform to a true cloud-ready infrastructure.

Running business-critical applications

The PowerOne is designed to provide the scalability, operations automation, security, and richness of services that business-critical applications require.

Whether it is a core banking function, the most critical ERP system, or the company billing platform, PowerOne offers the availability and performance requirements to match any business need. With the latest Intel CPUs in the densest compute chassis, the highest performing all-flash and NVMe-powered storage arrays, the leading virtualization platform, and the latest networking technology from Dell EMC, PowerOne is the best candidate to run any business-critical application.

Application-specific use case

The flexible configuration options of PowerOne allow you to build the specific platform setup you need to match an application or market's vertical requirements.

PowerOne can provide an efficient, highly available, and high-performing platform for:

- Healthcare (Epic EHR)
- Financial Services
- Oil and Gas
- Manufacturing
- Education

PowerOne sample use cases

- Service Providers

PowerOne components

This chapter describes the compute, network, and storage components of PowerOne.

PowerOne compute resources

This section describes the compute resources of PowerOne.

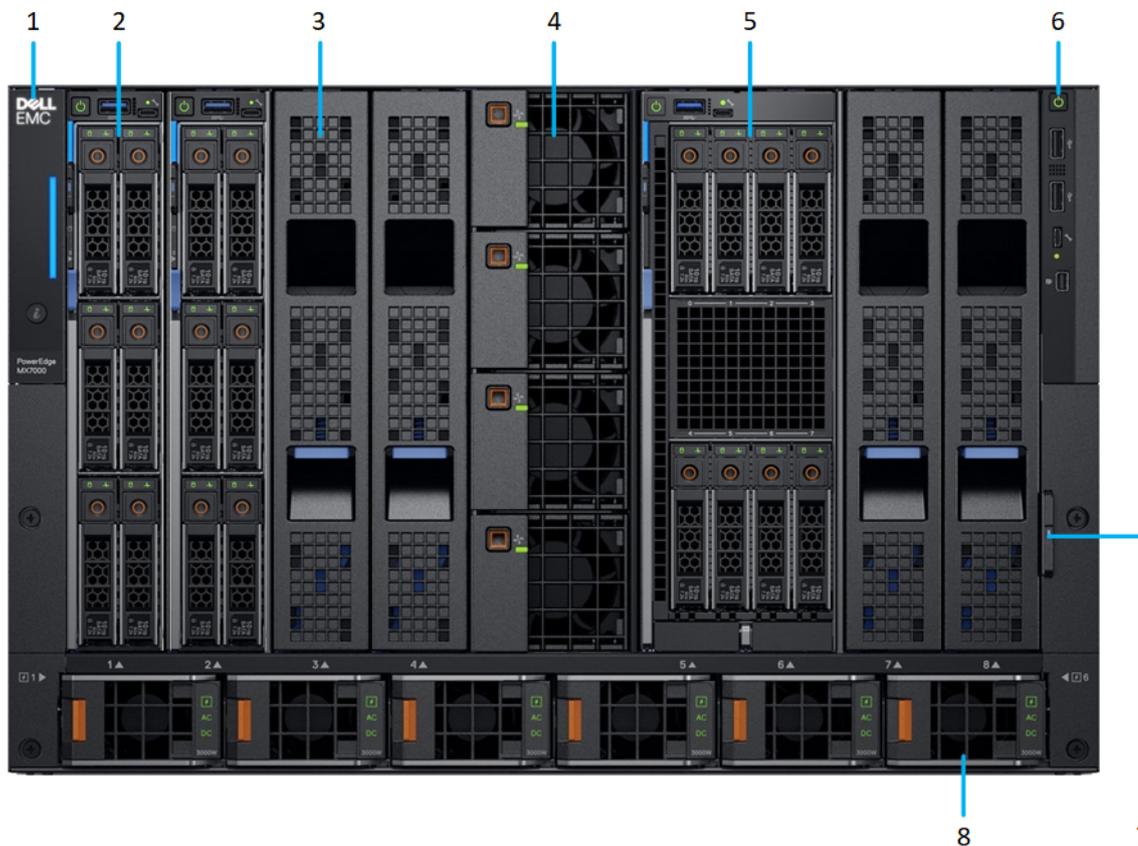
PowerEdge MX7000 chassis

PowerOne uses the Dell EMC PowerEdge MX infrastructure.

The PowerEdge MX7000 is a unified, high-performance infrastructure that provides agility, resiliency, and efficiency for virtualization workloads. The PowerEdge MX configuration used in PowerOne consists of a chassis infrastructure, compute sleds, and fabric switches.

Chassis front view

The following figure shows the front view of PowerEdge MX7000 in a PowerOne configuration.

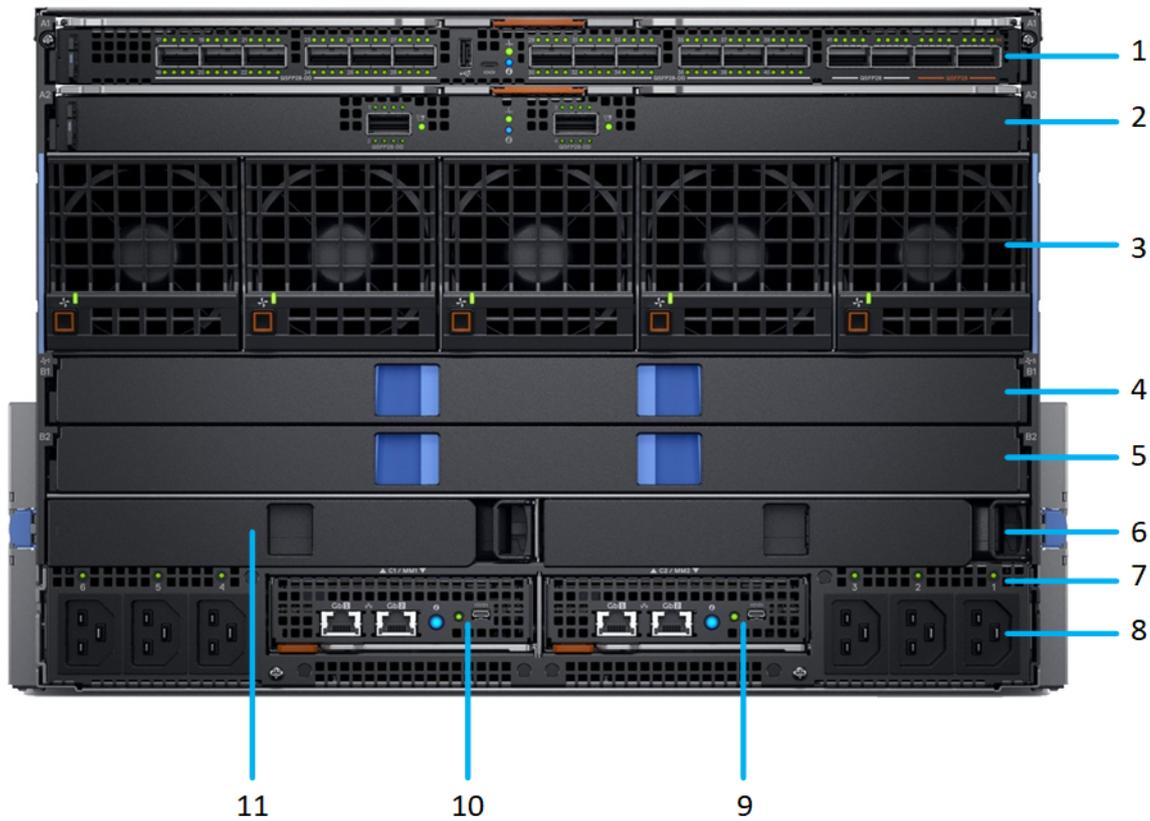


- | |
|---|
| 1. Left control panel |
| 2. Single-width compute sled (PowerEdge MX740c) |
| 3. Sled blank |

4. Front fan (4)
5. Double-width compute sled (PowerEdge MX840c)
6. Right control panel
7. Information tag
8. Power supply unit (8)

Chassis rear view

The following figure shows a rear view of the PowerEdge MX7000 in a PowerOne configuration.



1. IOM slot for Fabric A1 (PowerEdge MX9116n/MX7116n)
2. IOM slot for Fabric A2 (PowerEdge MX7116n/MX9116n)
3. Rear fans (5)
4. IOM slot for Fabric B1 (blank)
5. IOM slot for Fabric B2 (blank)
6. IOM slot for Fabric C2 (blank)
7. Power cord connection status LED
8. Power supply input sockets (6)
9. Management module 2
10. Management module 1

11. Slot for Fabric C1 (blank)

See the [Dell PowerEdge MX7000 Technical Guide](#) for additional information about the PowerEdge MX7000 chassis.

PowerEdge MX7000 power supply units

The PowerEdge MX7000 chassis that are part of a PowerOne compute domain are fully populated with six power supply units (PSUs) and configured in a grid- redundant configuration.

The PSUs are connected such that PSUs one through three (red) are associated with one grid, while PSUs four through six (blue) are associated with the second grid. The power sockets on the rear of the enclosure are set up based on color with three plugs on each side of the chassis for connectivity to the rack power distribution units (PDUs).

Note: A grid refers to a power distribution failure domain. Data centers have varying configurations of power distribution architectures.

PowerEdge MX7000 compute sleds

The PowerEdge MX7000 supports up to eight standard size (MX740c) or up to four double wide (MX840c) compute sleds.

PowerEdge MX740c compute sled

A fully featured, flexible, two-socket compute sled makes the PowerEdge MX740c ideal for virtualization.

The MX740c is a single width compute sled supporting up to 3 TB of RAM with one PCIe networking adapter. It uses the Boot Optimized Storage Solution (BOSS) for hypervisor installation. VMware ESXi is the only hypervisor with automatic provisioning capabilities in PowerOne Navigator.

See [Dell EMC PowerEdge MX740c Technical Guide](#) for additional information about the MX740c compute sled

Related information

[Appendix A - PowerEdge MX740c compute sled on page 50](#)

PowerEdge MX840c compute sled

A full-featured, flexible, four-socket compute sled, makes the PowerEdge MX840c ideal for virtualization.

It is a double width compute sled supporting up to 6 TB of RAM with two PCIe converged adapters. The MX840c uses the Boot Optimized Storage Solution (BOSS) for hypervisor installation. VMware ESXi is the only operating system with automatic provisioning capabilities in PowerOne Navigator.

See [Dell EMC PowerEdge MX840c Technical Guide](#) for additional information about the MX840c compute sled.

Related information

[Appendix B - PowerEdge MX840c compute sled on page 54](#)

Compute minimums and maximums

The following table lists the minimum and maximum compute resources that are supported in a single PowerOne:

	Minimum	Maximum	Notes
I/O Module (IOM) A1	1	1	<ul style="list-style-type: none"> Chassis 1 IOM A1 contains MX9116n Chassis 2 IOM A1 contains MX7116n All other chassis have MX7116n installed.

	Minimum	Maximum	Notes
I/O Module (IOM) A2	1	1	<ul style="list-style-type: none"> Chassis 1 IOM A1 contains MX9116n Chassis 2 IOM A1 contains MX7116n All other chassis have MX7116n installed
I/O Module (IOM) B1	0	0	Fabric B modules are not used in the PowerOne architecture. The addition of Fabric B modules may result in the chassis and all compute nodes in the chassis being placed in the component state <code>Invalid Configuration</code> .
I/O Module (IOM) B2	0	0	Fabric B modules are not used in the PowerOne architecture. The addition of Fabric B modules may result in the chassis and all compute nodes in the chassis being placed in the component state <code>Invalid Configuration</code> .
I/O Module (IOM) C1	0	0	Fabric C modules are not used in the PowerOne architecture. The addition of Fabric C modules may result in the chassis and all compute nodes in the chassis being placed in the component state <code>Invalid Configuration</code> .
I/O Module (IOM) C2	0	0	Fabric C modules are not used in the PowerOne architecture. The addition of Fabric C modules may result in the chassis and all compute nodes in the chassis being placed in the component state <code>Invalid Configuration</code> .
MX740c sleds per chassis	—	8	
MX840c sleds per chassis	—	4	
Double Wide sleds in slots	1-2, 3-4, 5-6, 7-8		
MX740c processors	1	2	Intel™ Xeon™ Scalable Processors, up to 28 cores per processor Intel® C628 Chipset with optional Intel® QuickAssist Technology
MX740c memory DIMMs	24	24	
MX740c memory capacity	8 GB (RDIMM) 64 GB (LRDIMM)	768 GB (RDIMM) 3 TB (LRDIMM)	See the Memory Specifications section of the Dell EMC PowerEdge MX740c Installation and Service Manual for details.
MX840c processors	1	4	Intel™ Xeon™ Scalable Processors, up to 28 cores per processor Intel® C628 Chipset with optional Intel® QuickAssist Technology
MX840c memory DIMMs	48	48	
MX840c memory capacity	8 GB (RDIMM) 64 GB (LRDIMM)	1.53 TB (RDIMM) 6.14 TB (LRDIMM)	See the Memory Specifications section of the Dell EMC PowerEdge MX840c Installation and Service Manual for details.

PowerOne networking resources

PowerOne brings together the following networking components and technologies:

- PowerEdge MX7000 networking components:
 - MX9116n Fabric Switching Engine (FSE)
 - MX7116n Fabric Expander Module (FEM)
- Dell S5232F-ON leaf switches that provide an aggregation point for the MX7000 chassis in a Pod. The S5232F-ON switches can provide an aggregation point for management connectivity when scaling beyond a single Pod.
- Dell S4148T-ON switches that provide management connectivity for a Pod.
- The Connectrix SAN switches that provide storage connectivity.

Dell EMC PowerEdge MX7000 networking

The Dell EMC PowerEdge MX platform is a next-generation modular platform that brings together compute, storage, and networking under a single management umbrella.

The PowerEdge MX7000 platform supports the following networking modules:

- Dell EMC MX9116n Fabric Switching Engine
- Dell EMC MX7116n Fabric Expander Module

PowerOne Internal Architecture

This section describes the internal architecture of PowerOne.

MX scalable fabric architecture

The MX scalable fabric architecture ties multiple chassis into a single network domain to behave like a single chassis from a networking perspective. It can scale up to 10 chassis, 80 compute sleds.

A scalable fabric consists of two main components — a pair of MX9116n Fabric Switching Engines (FSE) and MX7116n Fabric Expander Modules (FEM) used to connect remote chassis to the FSEs.

See [Dell EMC PowerEdge MX7000 Technical Guide](#) for more information about MX7000 networking.

Internal Fabrics

The PowerEdge MX7000 infrastructure has numerous communication paths or fabrics interconnecting the modules. The management fabric connects from each management module (MM) slot with a point-to-point link to each front slot and each IOM slot. This fabric runs at 1 Gb/sec. A 1 Gb/sec link spans the two MM slots for communication between these modules.

Input output fabrics

The internal connection of the main I/O subsystem in a PowerEdge MX7000 enclosure is referred to as Fabric A. This fabric is connected using a direct orthogonal connection between sleds in front of the chassis and I/O fabric modules in the rear. The mating connection on the sleds is on the mezzanine cards.

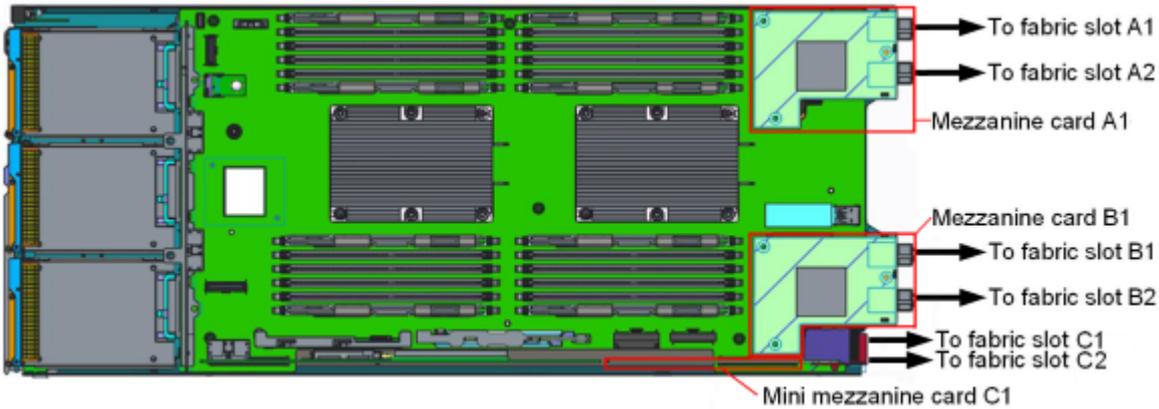
The fabric enables two IOMs (A1 and A2) for redundancy. Each IOM connects to all eight compute nodes through the orthogonal connectors. The connection from the IOM to each compute node (the mezzanine card on the compute node) is an x8 bi-directional link. Depending on the interface, all the lanes within the link may not be used; however the x8 provides a path for future high-bandwidth fabric types.

Front sled to rear port mappings

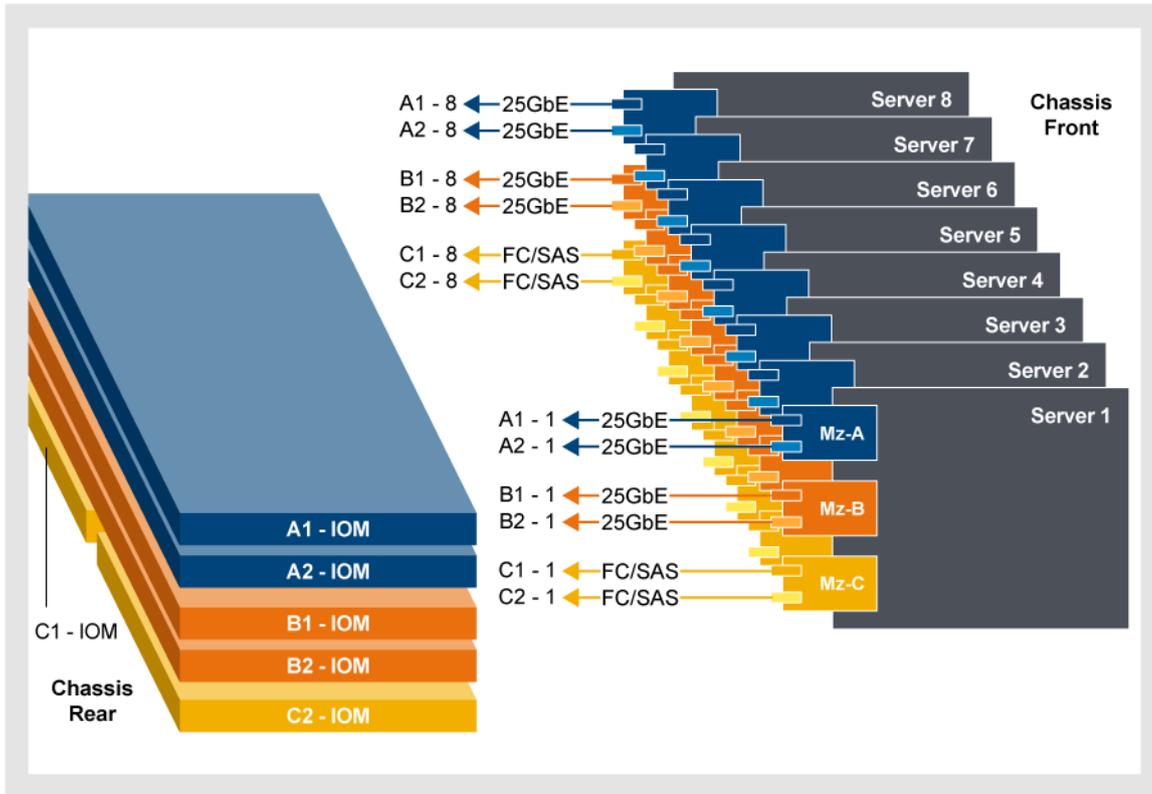
The external port assignments of pass-through modules are fixed due to the direct routing connection internal to the system and the IOM itself. A standard size compute sled that is configured with dual port Fabric A mezzanine maps to a

single port on each I/O module. Each compute sled number matches the IOM numbering on both the top and bottom I/O modules.

The following figure is an example of the MX740c mezzanine card slots.



The following figure shows all the compute sleds connecting to IOMs.

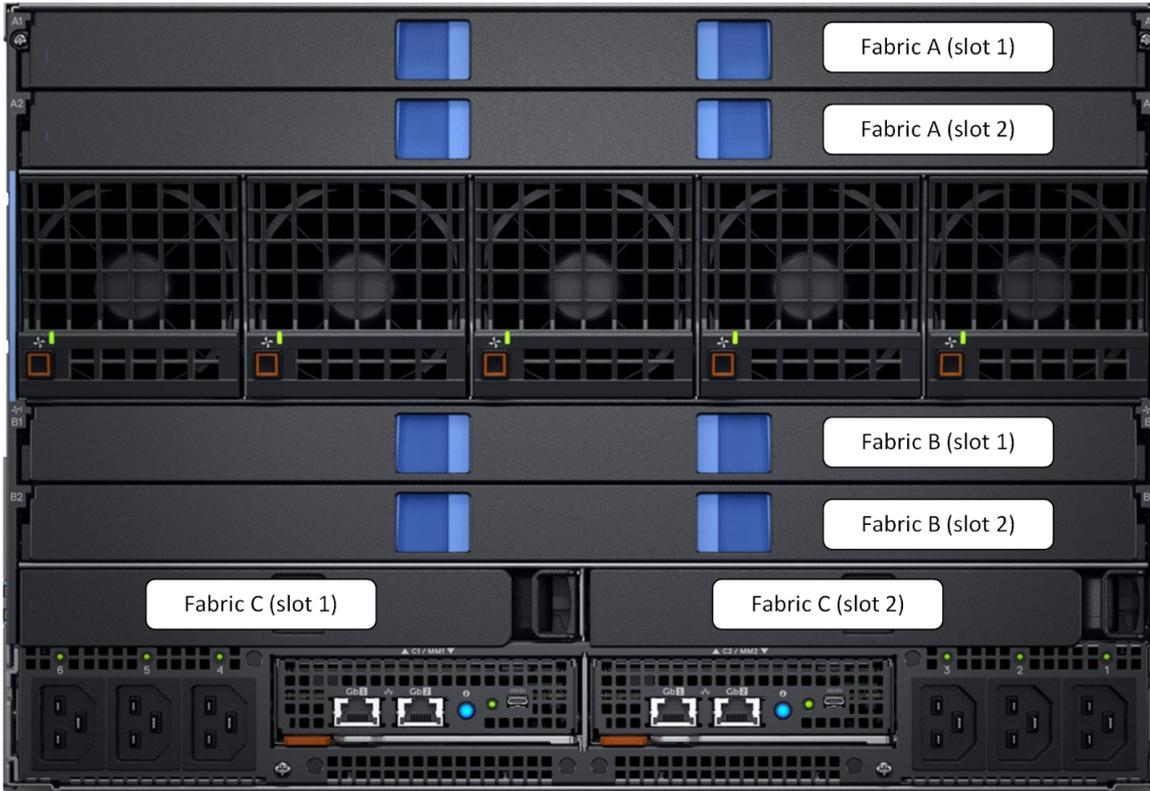


I/O modules

The IOMs A and B span the entire width of the chassis using direct orthogonal connections to the compute sleds. When installing these, the IOM will engage simultaneously with the Vertical Power Distribution Board (PDB) and the sleds.

PowerEdge MX7000 chassis

The following figure shows PowerEdge MX7000 chassis with two scalable networking fabrics (Fabric A and Fabric B) and one redundant storage fabric (Fabric C). In the current PowerOne architecture, only Fabric A is used. Fabric B and Fabric C are not used.



MX7000 networking modules

The MX7000 I/O modules used are MX9116n Fabric Switching Engine (FSE) and MX7116n Fabric Extender Module (FEM).

MX9116n Fabric Switching Engine

The MX9116n Fabric Switching Engine (FSE) is a scalable, high-performance, low latency 25-Gbps Ethernet switch purpose-built for the PowerEdge MX platform for demanding compute and storage traffic environments.

Each PowerOne compute domain can use either two or four MX9116n IOMs. These MX9116 IOMs are installed in the first two chassis of a compute domain and provide the interconnection point between the LAN (leaf and spine), SAN, and up to ten MX7000 chassis.



- 1 12 x QSFP28-DD
- 2 2 x QSFP28
- 3 2 X QSFP28 w/FC Capabiity

Port Speeds	
	100Gbps
	200Gbps

MX7116n Fabric Extender Module

The MX7116n architecture is extremely simple, consisting of a simple Ethernet repeater along with control circuitry for communication of low-level information (such as environmental information) to the MX7000 enclosure controller. Power and cooling is supplied by the MX7000 chassis.

PowerOne system fabric

PowerOne system fabric begins with a pair of Dell S5232F-ON (leaf) switches to support a single PowerOne Pod. The system fabric is designed for significant future scalability. In this design, an additional pair of S5232F-ON (leaf) switches will support each additional Pod. These S5232F-ON leaf switches will then connect to Dell Z9264F-ON (spine) switches.

For more information about these switches, see the following documents:

- https://i.dell.com/sites/csdocuments/Product_Docs/en/Dell-EMC-Networking-s5200-series-specsheet.pdf
- https://i.dell.com/sites/csdocuments/Product_Docs/en/dell-emc-networking-z9264-on-specsheet.pdf

PowerOne OOB management network

The Dell EMC PowerOne OOB management network consists of a pair of Dell S4148T-ON switches to support a single Dell EMC PowerOne Pod. The OOB network is designed for significant future scalability. In this design, the OOB management network will add a pair of Dell S4148T-ON switches to support each additional Pod. These S4148T-ON switches will then connect to Dell S5232F-ON (aggregation) switches.

For more information about these switches, see the following links:

- <https://i.dell.com/sites/doccontent/shared-content/data-sheets/en/Documents/dell-emc-networking-S4100-series-spec-sheet.pdf>
- https://i.dell.com/sites/csdocuments/Product_Docs/en/Dell-EMC-Networking-s5200-series-specsheet.pdf

PowerOne storage area network

The PowerOne SAN enables you to start with a small configuration, and then scale out to a larger configuration. The current per-Pod SAN configuration consists of two Connectrix MDS switches (9148T or 9396T).

The PowerOne SAN is designed for significant future scalability, to a core-edge design. In this design, the existing pair of SAN switches become edge switches, and a pair of Connectrix MDS switches (9396T, 9706 or 9710) are added to become the core. The edge will be expanded by adding a pair of edge switches for each additional PowerOne Pod. The core will be expanded by adding up to six additional core switches.

For more information about these switches, see the following link:

- https://www.dellemc.com/resources/en-us/asset/data-sheets/products/networking/h17267_connectrix_mds_9000t_switches_ss.pdf

PowerOne storage resources

PowerOne is engineered and designed to host a wide variety of business workloads. These workloads can have different storage requirements, in terms of performance, availability, or required data services.

PowerOne is designed to use the Dell EMC PowerMax 2000E/2000P/8000E/8000P external storage arrays. Local storage on the compute sleds is used only for Boot Optimized Storage Solution (BOSS) for the compute sleds' operating system.

PowerMax

The PowerMax family is the first Dell EMC hardware platform with storage back-end based in NVMe.

The NVMe-based PowerMax is created to provide performance benefits such as high bandwidth and IOPS, and low latency. NVM media offers high performance of host-based applications.

There are two primary hardware platforms in PowerMax: PowerMax 2000 and PowerMax 8000. Both use the modular Brick building concept, similar to the one used for the VMAX All Flash family.

See [Dell EMC PowerMax Family Overview](#) white paper for detailed configuration information.

Benefits of PowerMax

The primary benefits that PowerMax platforms offer are:

- An end-to-end NVMe storage design providing industry-leading IOPS density in a system with a small footprint.
- Next generation data storage media such as Storage Class Memory (SCM) and NVMe over Fabric (NVMe-oF) infrastructure.
- Efficient workload consolidation for block and file in a single platform.
- Best-in-class data services including data reduction using inline deduplication and compression.
- Ultrahigh reliability and serviceability for the enterprise.

PowerMax family

PowerOne supports PowerMax 2000 and PowerMax 8000. The following table provides the details of these two models.

PowerMax model	Benefits	Details
PowerMax 2000 (both 2000E and 2000P)	Provides efficiency and maximum flexibility in a 20U form factor.	<ul style="list-style-type: none"> • 1.7 M IOPs (8K RRH) • 64 FC/iSCSI ports • 1 to 2 Bricks per system • Up to one PBe per system of NVMe storage • All in a half rack (20U) • Open systems workloads only
PowerMax 8000 (both 8000E and 8000P)	Provides massive scale, performance, and IOPS density all within a two-floor tile footprint.	<ul style="list-style-type: none"> • 10 M IOPs (8K RRH) • 256 FC/FICON/iSCSI ports • 1 to 8 Bricks per system • Up to 4 PBe per system of NVMe storage • All in two racks (two-floor tiles) • Open systems/mainframe/mixed workloads

Both PowerMax arrays use Dynamic Virtual Matrix architecture and a new version of HYPERMAX OS rewritten for the NVMe platform called PowerMax OS 5978. PowerMax OS can run natively on both PowerMax systems and on legacy all-flash systems as an upgrade. PowerMax provides trusted data services, with improved simplicity, capacity, and performance for highly virtualized environments demand, while meeting the economic needs of the more traditional storage workloads. In addition, PowerMax supports deploying applications such as real-time analytics, machine learning, and big data that demand lower storage latency and higher IOPS densities.

PowerMax architecture

PowerMax configurations consist of modular building blocks called PowerMax Bricks (Bricks). The modular Brick architecture reduces complexity and enables easier system configuration and deployment. This architecture also enables the systems to scale while continuing to deliver predictable high performance.

For PowerOne, the open systems Brick supports Fibre Channel configurations. The Brick can also be configured for file storage using embedded NAS.

The initial Brick includes a single engine consisting of two directors, two system power supplies (SPS), and two 24-slot 2.5" NVMe Drive Array Enclosures (DAE24) preconfigured with an initial total usable capacity. The initial capacity depends on the type of PowerMax system and RAID configuration. Each Brick comes preloaded with PowerMax. The Brick concept enables PowerMax to scale up and scale out. Scaling up is achieved by adding Flash Capacity Packs, that

add a certain usable storage capacity that depends on the RAID protection type selected. Scaling out is achieved by adding additional Bricks to produce predictable linear performance improvements, regardless of the workload.

See the [Dell EMC PowerMax Family Overview](#) for more information.

PowerMax Engines

The core of the Brick is the engine. It is the central I/O processing unit, which is built with redundancy for high availability.

Each Brick consists of:

- Redundant directors that contain multicore CPUs and memory modules
- Interfaces to universal I/O modules, such as, back-end, front-end, InfiniBand, and flash I/O modules

The communication backbone of the Brick is the Trusted Dynamic Matrix architecture. The virtual matrix enables interdirector communications over redundant internal InfiniBand fabrics. The InfiniBand fabric provides a foundation for a highly scalable, extremely low latency, and high-bandwidth backbone which is essential for an all-flash array. This capability is also essential to enable PowerMax to scale upwards and scale outwards in the manner that it does.

Each engine has two directors. Each director comprises dual CPU sockets that can support multi-core multithreaded Intel processors.

PowerMax model	Engine CPU type	CPU cores	Cores per Brick engine	Max cores per system
PowerMax 2000	Dual Intel Broadwell, 12 core, 2.5 GHz	24	48	96 (2 bricks maximum)
PowerMax 8000	Dual Intel Broadwell, 18 core, 2.8 GHz	36	72	576 (8 bricks maximum)

The Brick engine uses a core pooling mechanism that can dynamically load-balance the cores by distributing them to the front end, back end, and data services (such as SRDF, eNAS, and embedded management) running on the engine. The core pools can be tuned to shift the bias of the pools at any time to front-end heavy or back-end heavy workloads to further optimize the solution for a specific use case.

Every director has 16 memory slots that can be populated with 32 GB and 64-GB DDR4 DIMMS to achieve up to 1-TB cache per director (2-TB cache maximum per Brick engine).

PowerMax model	Cache per Brick	Max cache per system
PowerMax 2000	512 GB, 1 TB, 2 TB	4 TB (2 Bricks maximum)
PowerMax 8000	1 TB or 2 TB	16 TB (8 Bricks maximum)

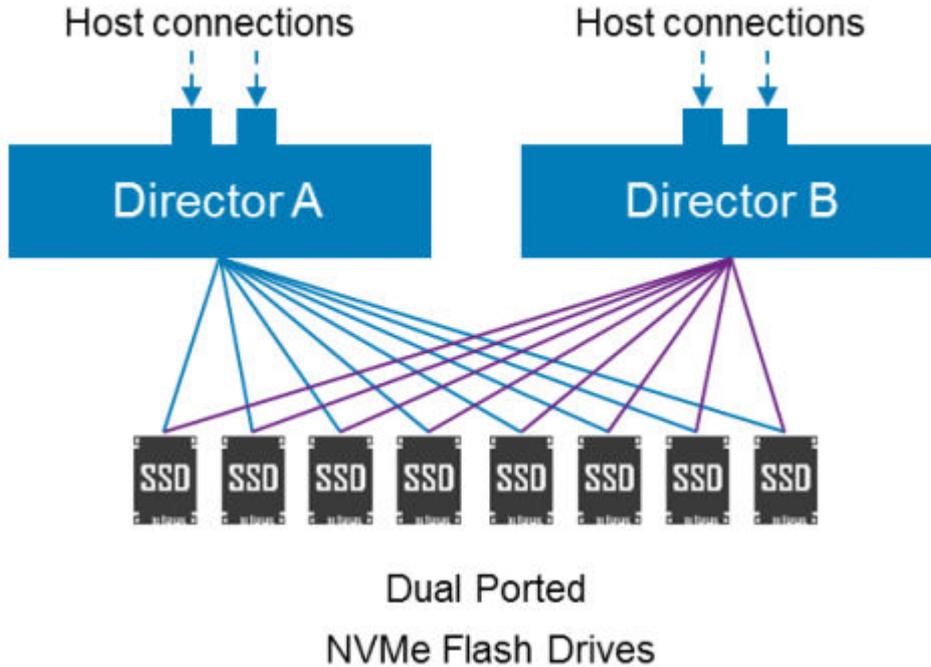
PowerMaxOS

PowerMaxOS continues to provide industry-leading high availability, I/O management, quality of service, data integrity validation, data movement, and data security within an open application platform. PowerMaxOS uses a real-time, nondisruptive storage hypervisor that manages and protects embedded services by extending high availability to services that traditionally would have run external to the array. The primary function of PowerMaxOS is to manage the core operations performed on the array:

- Processing I/O from hosts
- Implementing RAID protection
- Optimizing performance by enabling direct access to hardware resources
- Managing and monitoring the system

PowerMax Smart RAID

PowerMax uses a new active/active RAID group accessing scheme called Smart RAID that enables RAID groups to be shared across directors, providing each director active access to all drives on the Brick.



The use of Smart RAID on PowerMax provides customers with performance benefits, as both directors on an engine can drive I/O to all the flash drives. This creates balanced configurations in the system regardless of the number of RAID groups.

PowerOne network

The PowerOne network provides three networks that facilitate connectivity for the entire PowerOne. The configuration of these networks is automated by the PowerOne Controller. The following table lists these networks.

Network	Description
System fabric	Provides connectivity for east-west production traffic between PowerOne hosts and north-south production traffic between these hosts and production network of the customer.
Management network	Provides OOB management connectivity between all PowerOne components and management network of the customer.
Storage Area Network(SAN)	Provides connectivity for all the storage traffic between external storage and hosts within PowerOne.

System fabric

The system fabric is a key component of PowerOne. It provides high speed with low latency throughput and enables:

- Physical scaling from small to very large network configurations.
- Automated network configuration, facilitated by the PowerOne Controller.

The system fabric consists of the following physical network components:

- Dell EMC Networking S5232F leaf switches
- Dell EMC MX7000 Networking IO Modules (IOM):
 - Dell EMC PowerEdge MX9116 Fabric Switching Engine (FSE)
 - Dell EMC PowerEdge MX7116 Fabric Expanding Module (FEM)

Fabric connectivity

PowerOne System Fabric enables you to start with a small configuration and scale out to a large configuration. The current per-Pod configuration is for two Dell EMC Networking S5232F-ON switches.

PowerOne is designed for significant future scalability: by adding additional S5232F-ON switches (in pairs) and connecting them through Dell EMC Networking Z9264F-ON (spine) switches. Maximum scalability potential includes six Z9264F-ON (spine) switches and 64 S5232F-ON (leaf) switches.

Each pair of S5232F-ON switches provides network connectivity for a single PowerOne Pod. The following figure shows the port layout for each S5232F-ON switch. The port layout shows the following ports:

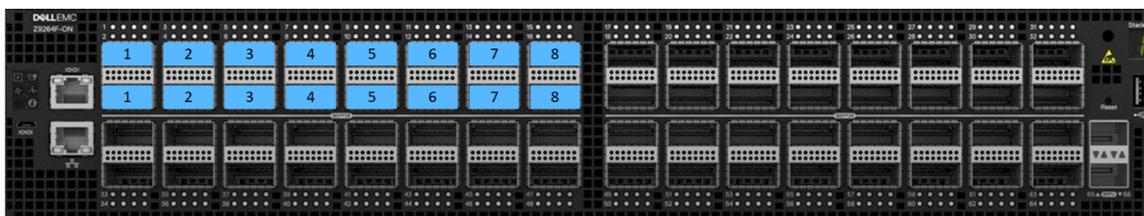
- 12 ports reserved for connections to the compute domains
- Two ports used for an Interswitch Link (ISL) between the two S5232F-ON switches
- Six ports reserved for connections to the spine switches for future scalability



1 Ports 1-12 (Connection to compute pod) 3 Ports 25-30 (Connection to spine switches) Port speeds: 100Gbps

2 Ports 23-24 (ISL)

PowerOne is designed for significant future scalability, by incorporating Z9264F-ON (spine) switches, resulting in a leaf-spine architecture. To support this scenario, the port layout in the following figure shows one physical connection (100 GbE) between each S5232F-ON (leaf) switch and each Z9264F-ON (spine) switch for the eight Pods.



Port speeds: 100Gbps

1 Ports 1-2 (Pod 1) 3 Ports 5-6 (Pod 3) 5 Ports 9-10 (Pod 5) 7 Ports 13-14 (Pod 7)

2 Ports 3-4 (Pod 2) 4 Ports 7-8 (Pod 4) 6 Ports 11-12 (Pod 6) 8 Ports 15-16 (Pod 8)

All the highlighted ports in blue in the figures above are of form factor QSFP28, which provides four channels at 28 Gbps each. By driving these channels at different speeds and/or combining channels, various Ethernet speeds can be achieved. In this case, all four lanes are combined into a single instance, resulting in a 100 Gb connection.

The following cables types are supported:

- DAC-Q28-100G (Direct Attach Copper)
- AOC-Q28-100G (Active Optical Cable)

Pod connectivity

Each pair of S5232F-ON switches provides network connectivity for a single Pod. A Pod supports up to three compute domains, which in turn support up to 10 PowerEdge MX7000 chassis (up to 80 compute sleds). Therefore, each Pod can support up to 30 MX7000 chassis (up to 240 blade servers).

Only the first two chassis in each compute domain connect directly to the S5232F-ON switches. These connections are made from the MX9116n IOMs. The MX7116n IOMs connect only to the MX9116n IOMs.

The following table shows the details of the IOM configuration supported by PowerOne.

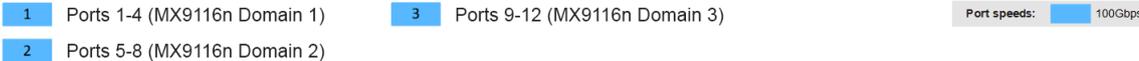
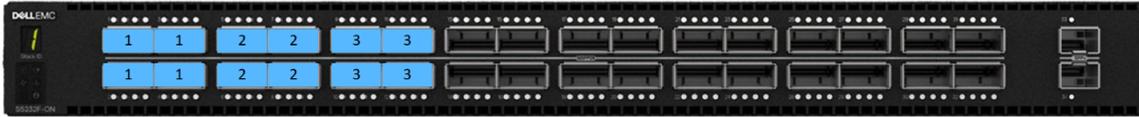
	Fabric A	
Chassis	Slot A1	Slot A2
1	MX9116n	MX7116n
2	MX7116n	MX9116n
3-10	MX7116n	MX7116n

Domain connectivity

Each PowerEdge MX7000 compute domain must connect to a pair of Dell EMC Networking S5232F-ON leaf switches to enable traffic flow to and from the compute sleds.

Dell EMC Networking S5232F-ON switch to Dell EMC MX9116n Fabric Switching Engine

The Dell EMC S5232F-ON Switch connects directly to the first two chassis in each compute domain, using the MX9116n IOMs.



Dell EMC MX9116 Fabric Switching Engine IOM

On the MX9116n IOM, each of the first 12 external ports is of form factor QSFP28-DD, which provides two 100 GbE ports per socket, enabling a single connector to support 200 Gbps. In addition, the first 16 ports are internal ports, so faceplate numbering starts at port 17. The port layout consists of the following ports:

- 12 x QSFP-DD (ports 17 through 40)
- 2 x QSFP28 (ports 41 through 42)
- 2 x QSFP28 w/FC support (ports 43 through 44)

On each MX9116n IOM, four ports are used to connect to a pair of S5232F-ON switches. Two ports are connected to each S5232F-ON switch. In addition, four ports are used for an Interswitch Link (ISL) between the two MX9116n IOMs.

The MX9116n connection to the first S5232F-ON switch requires a breakout cable (2 x 100 Gbps) because the ports on the S5232F-ON switch are of the form factor QSFP28.

Port Connectivity

Ports 27–28 (QSFO28-DD) connect to QSFO28 ports on the first S5232F-ON switch and require a breakout cable. The following cables are supported:

- QSFO28DD-2Q28-100G (Direct Attach Copper)
- AOC-Q28DD-2Q28-100G (Active Optical Cable)

Ports 37–40 (QSFO28-DD) connect to the same ports on the other MX8216n IOM and so breakout cables are not required. The following cables are supported:

- DAC-Q28DD-200G (Direct Attach Copper)
- AOC-Q28DD-200G (Active Optical Cable)

Ports 41–42 (QSFO28-DD) connect to QSFO28 ports on the other S5232F-ON switch. The following cables are supported:

- DAC-Q28-100G (Direct Attach Copper)
- AOC-Q28-100G (Active Optical Cable)

The following diagram shows the port connectivity.



- 1 Ports 27-28 (S5232)
- 2 Ports 37-40 (ISL)
- 3 Ports 41-42 (S5232)

Port Speeds

- 100Gbps
- 200Gbps

Dell EMC MX9116n to Dell EMC MX7116n

Connections made from the MX7116n IOM to the MX9116n IOM are always in the same fabric. No connections are made between fabrics, ensuring the port configuration on each MX9116n is the same.

The following figure shows chassis connections that are positioned to provide a balanced airflow.



- | | |
|-----------------------------|----------------------------|
| 1 Ports 17-18 (Chassis 1/2) | 7 Ports 29-30 (Chassis 3) |
| 2 Ports 19-20 (Chassis 4) | 8 Ports 31-32 (Chassis 5) |
| 3 Ports 21-22 (Chassis 6) | 9 Ports 33-34 (Chassis 7) |
| 4 Ports 23-24 (Chassis 8) | 10 Ports 35-36 (Chassis 9) |
| 5 Ports 25-26 (Chassis 10) | |

Port Speeds

- 200Gbps

On the MX7116n IOM, only ports 1 through 2 are used and a single connection is made to an MX9116n IOM. Ports 3 through 4 are reserved for future use.



- 1 Port Group 1 (MX9116n)

Port Speeds

- 200Gbps

Port Connectivity

All port pairs are of form factor QSFP28-DD, which provides eight lanes at 25 Gbps each. By driving these channels at different speeds or combining channels, various Ethernet speeds can be achieved. In this case, the channels are used individually resulting in 8 x 25-GbE links. This enables each compute sled to have dedicated 25 GbE channels.

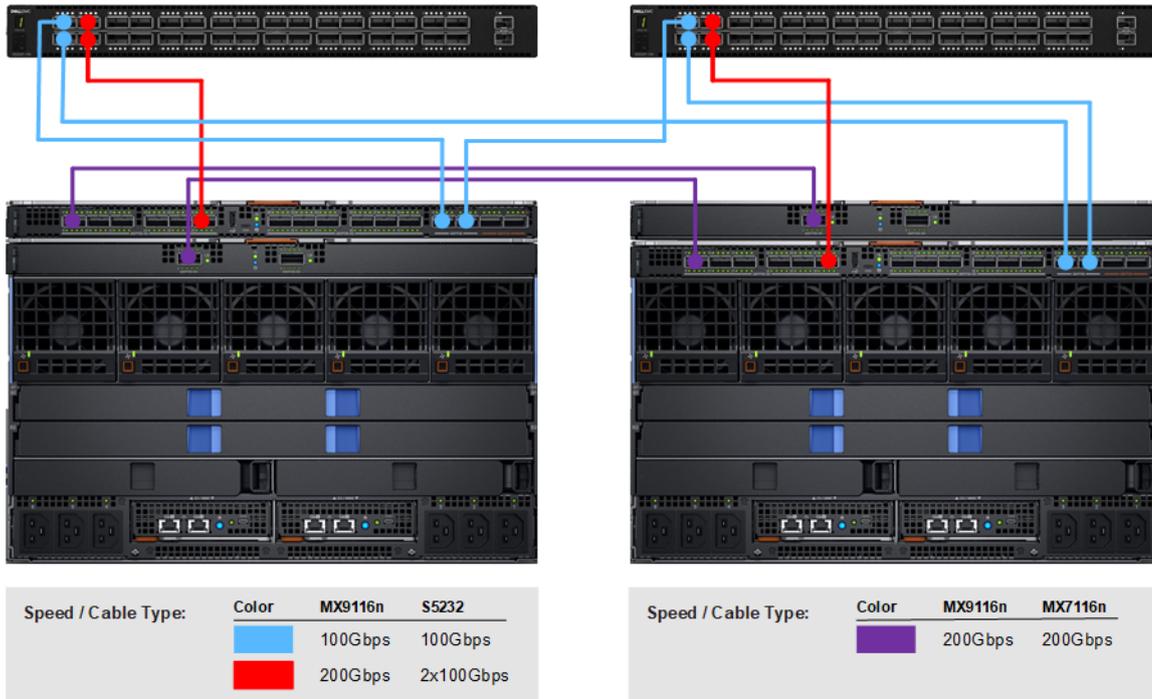
These ports connect to ports on another IOM, so breakout cables are not required. The following cables are supported:

- DAC-Q28DD-200G (Direct Attached Copper)
- AOC-Q28DD-200G (Active Optical Cable)

Two MX9116n configuration domain connectivity

The following figure shows domain connectivity for the MX9116n configuration (first two chassis only).

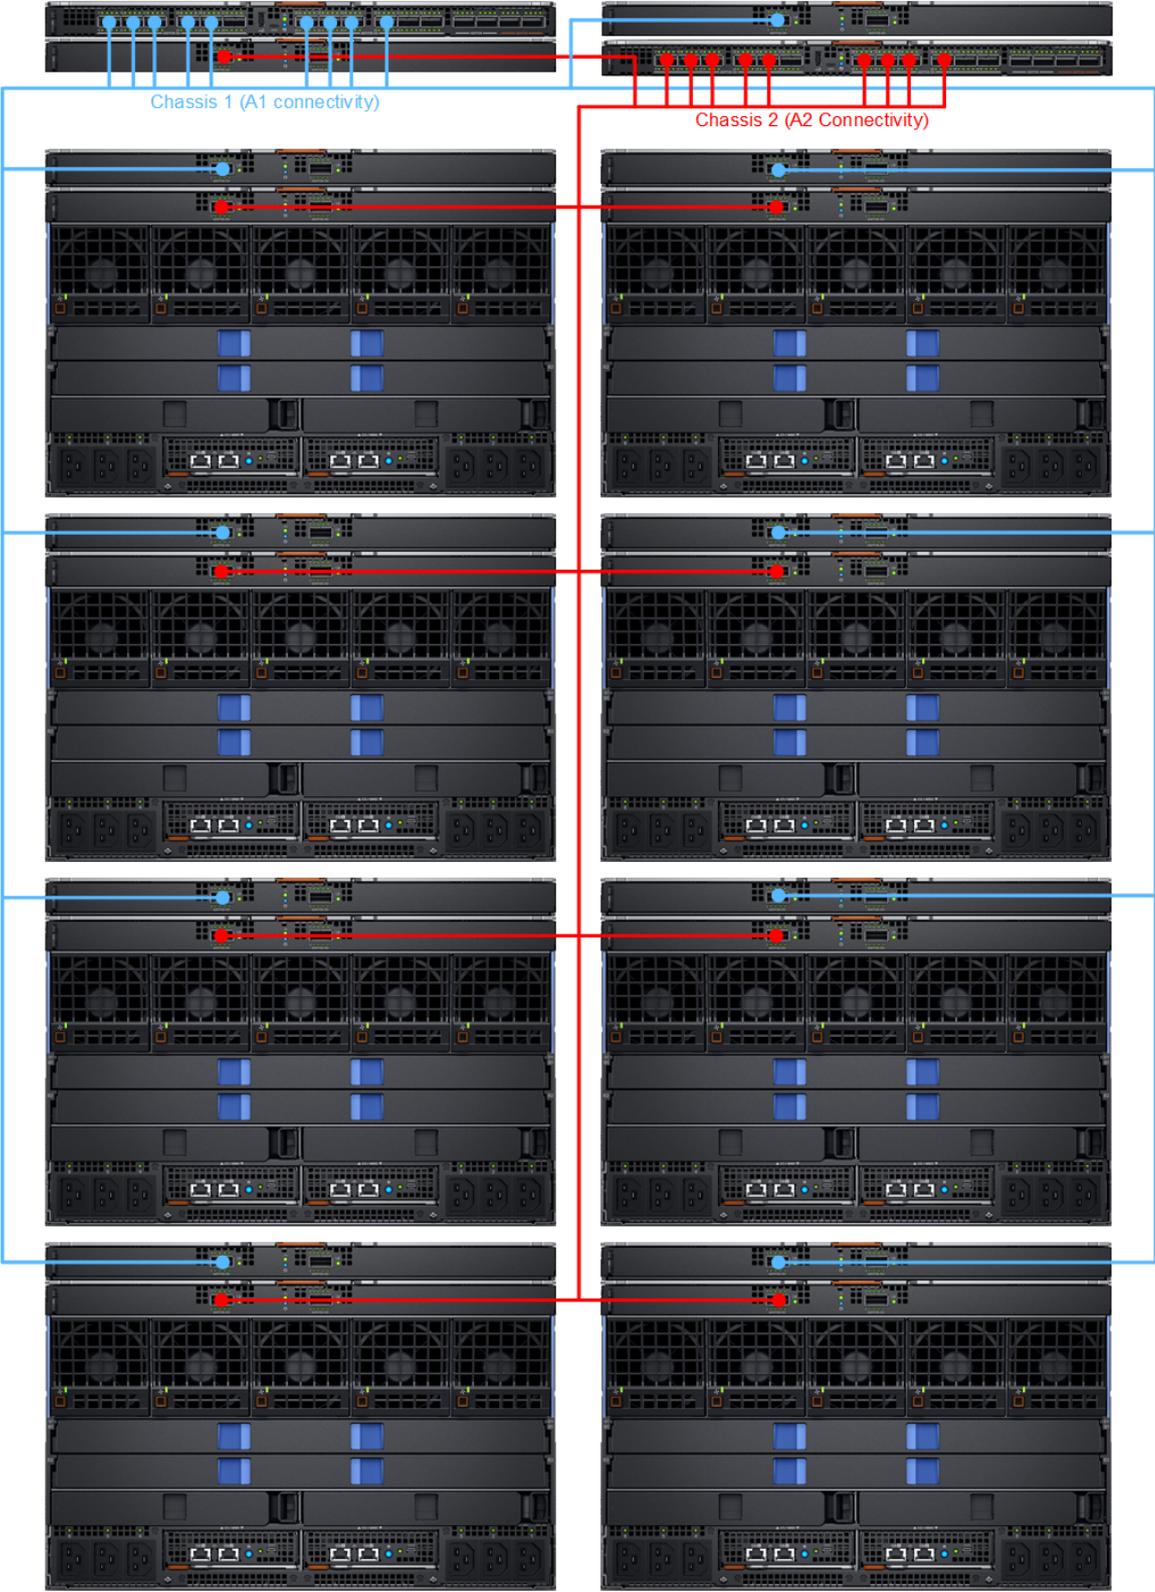
PowerOne network



MX7116n domain connectivity

The following figure shows domain connectivity for chassis 3–10.

Note: Only the MX9116n and MX7116n IOMs are shown for chassis 1 and 2.



PowerOne SAN fabric

Fibre Channel over Ethernet (FCoE) is used for storage I/O in PowerOne. Only Fabric A is used for FCoE. On Fabric A, the connection between the compute sleds and the MX9116n IOMs are configured for unified IO (network and storage I/O).

The MX9116n IOMs on Fabric A are directly connected to Connectrix SAN switches.

FC connections (Fabric A only)



1 Ports 43-44 (MDS Edge Switch)

Port speeds: 100Gbps

The MX9116n IOMs support the following two interoperability modes for Fibre Channel storage:

- NPIV Proxy Gateway (NPG) mode
- Direct Attached FC (F_Port) mode

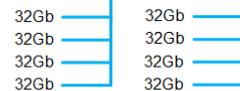
PowerOne uses the NPG mode along with the core-edge, and this enables a better approach to maintain performance and availability as the system scales out.

A breakout cable is used to connect each port on MX9116n to four 32 Gb ports on the MDS edge. Eight ports are used on the MDS edge switch.

FC connection breakout



Port speeds: 100Gbps



NPIV and oversubscription

N-Port ID Virtualization (NPIV) enables multiple N-ports to perform a Fabric Login (FLOGI) into a single F-port. Up to a maximum of 80 compute sleds can connect through these connections. Therefore, up to ten N-ports can be virtualized through each connection.

This feature simplifies the Fibre Channel cabling requirements and maximum F-port utilization. However, it adds an additional design consideration. Over subscription must be considered at each F-port and not just at the E-ports (ISL) now.

In PowerOne, a single compute domain can support up to eighty compute sleds. These compute sleds are connected to each SAN fabric through eight physical connections. Up to ten server blades can be connected to the storage through each physical connection. Each port on the 9396T or 9148T switches supports up to ten virtual N-ports.

Storage connectivity and scale-out

In a PowerOne with a single Pod, the storage arrays can be connected to the 9396T or 9148T switches.

9396T switch

The 9396T switch provides ninety-six (96) 32-GB FC ports.

For more information about the 9396T switch, refer to the following link:

https://www.dellemc.com/resources/en-us/asset/data-sheets/products/networking/h17267_connectrix_mds_9000t_switches_ss.pdf

Configuring SAN traffic

VSANs

The term virtual storage area network (VSAN) is used in the context of Fibre Channel switching and it suggests the segregation of SAN traffic into logical compartments. It does not suggest the VSAN feature that is provided by VMware.

Most PowerOne systems require a single VSAN (per SAN fabric). PowerOne supports multiple VSANs (per SAN fabric).

Zoning

PowerOne uses Smart Zoning to eliminate the need to create single initiator zones. Smart Zoning adds keywords that enable the designation of N-ports as initiator, target, or both.

The following guidelines are used to implement Smart Zoning on the PowerOne SAN switches:

- Use Smart Zoning as the default zoning method for segregated networking of the production VSANs.
- Enable Smart Zoning on the primary VSANs (VSAN 10 on Fabric A and VSAN 11 on Fabric B).
- Create device aliases for all components.

Device aliases

Note: Fabric-wide limit is 8000 and FC alias (VSAN-specific) limit is 2000.

PowerOne management network

The PowerOne management network provides OOB connectivity for PowerOne.

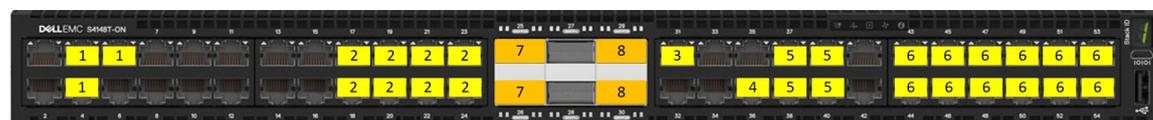
Similar to the PowerOne system fabric, the OOB network starts small and is designed for significant future scalability. The current OOB network consists of two Dell EMC S4148T-ON switches. When PowerOne is scaled out, an additional pair of S4148T-ON switches will be deployed for each Pod. To accommodate customer connectivity for these additional S4148T-ON switches, two S5232F-ON switches will be deployed as OOB aggregation switches.

Physical architecture

This section describes the OOB connections.

OOB connections (common to each Compute Pod)

The following diagrams show OOB management connections that are common to any Compute Pod.



- | | | | |
|---|---------------------------------------|---|---|
| 1 | Ports 3-5 (Compute Domain 1-3, MCM-1) | 6 | Ports 43-54 (Reserved for storage) |
| 2 | Ports 17-24 (Reserved for IPI) | 7 | Ports 25-26 (ISL) |
| 3 | Ports 31 (Leaf switch #1) | 8 | Ports 29-30 (Uplinks to aggregation switch) |
| 4 | Ports 36 (MDS 9148T #1) | | |
| 5 | Ports 37-40 (Reserved for IPI) | | |

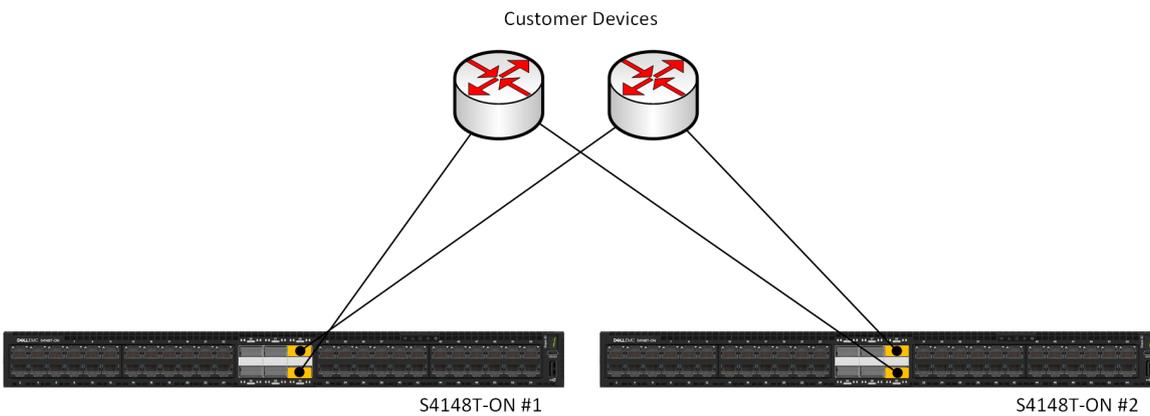
Port Speeds	
	10Gbps
	40Gbps



- | | | |
|--|---|--|
| <ul style="list-style-type: none"> 1 Ports 3-5 (Compute Domain 1-3, MCM-2) 2 Ports 17-24 (Reserved for IPI) 3 Ports 31 (Leaf switch #2) 4 Ports 36 (MDS 9148T #2) 5 Ports 37-40 (Reserved for IPI) | <ul style="list-style-type: none"> 6 Ports 43-54 (Reserved for storage) 7 Ports 25-26 (ISL) 8 Ports 29-30 (Uplinks to aggregation switch) | <p>Port Speeds</p> <ul style="list-style-type: none"> 10Gbps 40Gbps |
|--|---|--|

OOB customer connection

The following diagram shows the S4148T-ON OOB management switches that are directly linked to the management network of the customer environment.



Configuration of the management network

In this configuration, each uplink from the S4148T-ON switches is a point-to-point L3 connection. The OOB management networks are advertised into the customer’s management network through BGP or OSPF.

In PowerOne there is an L3 connection between the S4148T-ON (OOB) switches and the S5232F-ON (leaf) switches in the first Pod.

PowerOne platform organization

PowerOne supports the VMware platform.

Four compute sleds are reserved for the vSphere Management Cluster: the first two compute sleds in chassis 1 and the next two compute sleds in chassis 2.

The preferred VMware platform layout follows the guidance that is provided in the VMware Validated Design (VVD) documentation. For more information about VMware VVD, see the VMware Validated Design product documentation available at the following location:

<https://docs.vmware.com/en/VMware-Validated-Design/index.html>

VMware Validated Design (VVD) considerations

Management cluster

In the VVD architecture, one ESXi cluster is defined for host management and orchestration of virtual machines. In PowerOne, this ESXi cluster is deployed on the vSphere Management Cluster nodes. A dedicated instance of vCenter manages this ESXi cluster.

NSX-v

In VVD 5.0.1, NSX-v is used to provide connectivity to VMware management virtual machines. The VMware management network is different from the PowerOne management network. Connectivity from the customer environment to the VMware management is through the PowerOne System Fabric.

Compute/Edge cluster

In the VVD architecture, additional ESXi clusters are defined to host workload virtual machines. One of these ESXi clusters is defined as the compute/edge. In addition to hosting workload virtual machines, this ESXi cluster forms the edge of the VMware compute networks. PowerOne does not require compute edge cluster, but it can be deployed on a workload CRG.

NSX-T (optional)

In VVD 5.0.1, NSX-T is used to provide connectivity to VMware workload virtual machines. Connectivity from the customer environment to the VMware compute networks is through the PowerOne System Fabric.

PowerOne storage

This chapter describes the details of PowerOne storage.

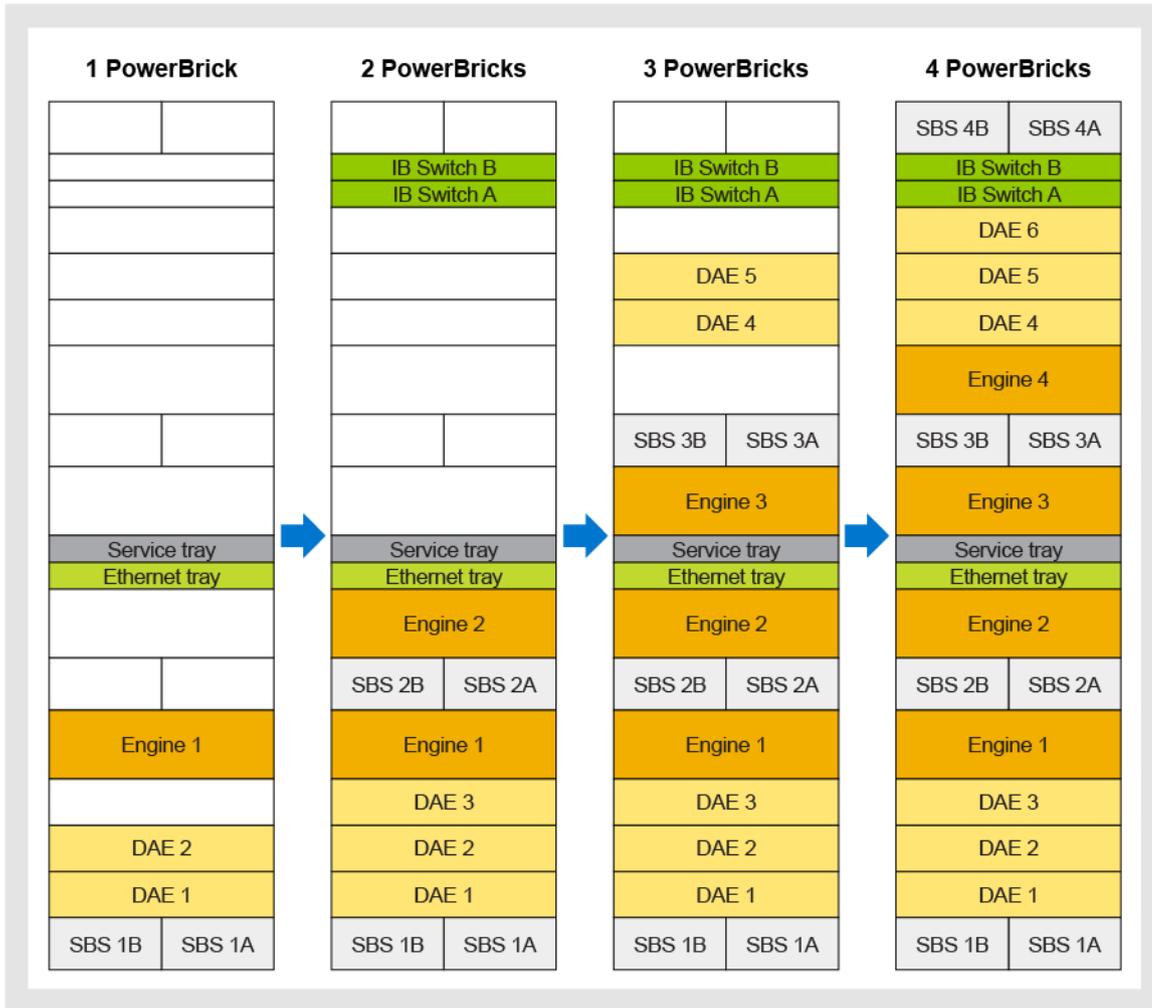
PowerMax

PowerOne incorporates the PowerMax 2000 and PowerMax 8000 storage arrays.

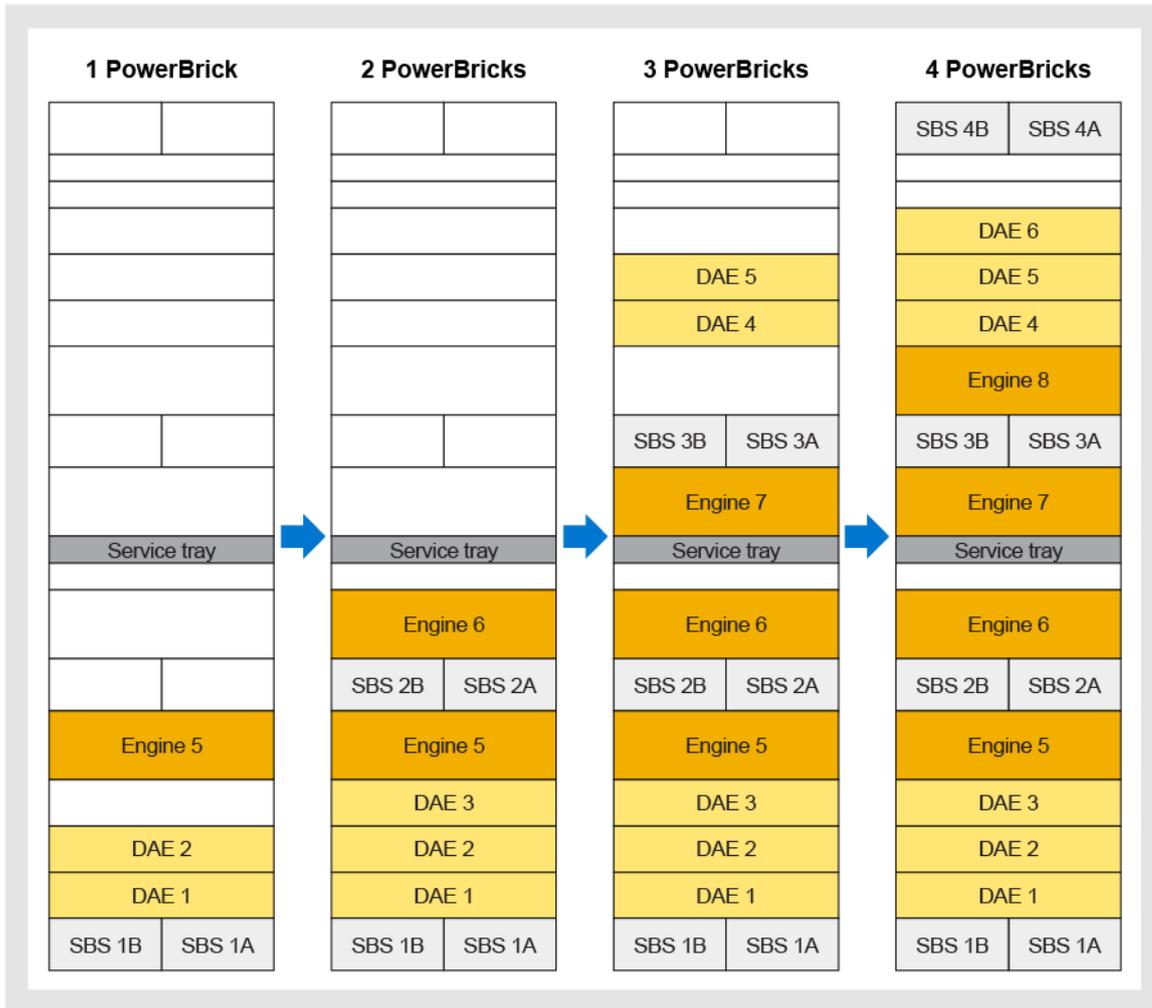
PowerMax 2000 and PowerMax 8000 are built based on Bricks. Each Brick consists of an engine, dual standby power supplies (SPS), and one or two NVMe drive array enclosures (DAE). Odd and even engines share a DAE. Dual Ethernet switches and InfiniBand (IB) switches interconnect to all Bricks. A slide out work tray provides space for a service laptop with power and management network access to the engines.

The following figures show how scaling affects the contents of the two racks, from one to eight Bricks in a PowerMax 8000.

Cabinet One



Cabinet Two



PowerMax Integration

As an enterprise solution, PowerOne implements integration points with all the system components.

PowerMax MMCS provides environmental monitoring capabilities for power, cooling, and connectivity. Each MMCS has two network connections that connect to the local area network (LAN) of the customer data center. One connection enables the monitoring of the system and remote connectivity for the Dell EMC Customer Support team, and the second connection for the Network Address Translation (NAT) Gateway.

On PowerMax 2000 and PowerMax 8000 models, there are two MMCS modules that are installed on two directors of each Brick. Within PowerOne, MMCS Ethernet ports on Brick-1 are connected to OOB management switches A and B as physical integration points for the PowerMax control plane.

PowerMax Management

PowerMax systems provide multiple element manager options. Each of the management options has different implementation methods and provides different capabilities. The PowerOne Controller Software is integrated with PowerMax at the element manager level.

Management installation options	Summary	Require gatekeeper	Unisphere UI	RestAPI	Dell EMC Solutions Enabler	Solutions Enabler Client/ Server
Embedded management (eMGMT)	PowerMax with HyperMax/OS introduces embedded management (Unisphere and Solutions Enabler (SE) with SMI software package) on a host container on the array (MMCS modules).	No	Yes	Yes	Yes. However, SE is available to Dell EMC personnel only.	Yes. SE client that is installed on VMs on the vSphere Management Cluster can point to the eMGMT SE server.
Solutions enabler client/server	PowerMax SE is installed on a virtual machine on the vSphere Management Cluster. The PowerMax SE client is configured as a remote client and points to the array host container eMGMT SE server.	No	No	No	Yes	Yes
PowerMax management virtual appliance	PowerMax appliance runs on one of the controller platform ESXi servers with PowerMax Gatekeeper. The virtual appliance provides Unisphere SE with SMI to manage the PowerMax System.	Yes	Yes	Yes	Yes	Yes

PowerMax storage provisioning standards

This section describes various supported standards associated with storage provisioning standards.

Storage resource pool (SRP) standards

PowerMax arrays are preconfigured from Dell EMC with one storage resource pool (SRP) containing all the storage available to the array.

Service level objectives (SLO) standards

During the preconfiguration process, all the NVMe drives are added to the disk groups based on the technology, capacity, and the required protection scheme. Each disk group has capacity that is partitioned for use by FBA devices, which are then added to the preconfigured SRP. The storage administrator can display the configuration of the SRP on the array and generate reports detailing the demand that storage groups or SLOs are placing on the SRP.

Fully Automated Storage Tiering (FAST) is controlled through predefined SLOs and Workloads.

Data reduction option standards

PowerMax delivers extreme efficiency with inline deduplication and compression, delivering up to 5:1 data reduction (average of 3:1), space efficient snaps, and thin provisioning. Its inline deduplication and compression have virtually zero impact on performance, and can be used with all data services, and are turned on/off by the application.

In a PowerMax system with Data Reduction hardware, by default, all storage groups have data reduction (compression and deduplication) enabled. This can be disabled based on the storage group from the element manager.

Data LUN standards

Data devices are thin devices (TDEV) of varying sizes depending on customer requirements.

By default, data devices have compression that is enabled at the storage group level. Disable compression if the application is not designed for compression.

The use of host LUN ID 0 is permitted for data LUNs when all hosts in the cluster are booting from local SD.

Storage masking standards

Storage is presented to the servers through a masking view.

Masking views consist of the following:

- **Initiator group**—collection of initiator WWPNs
- **Port group**—collection of array Front-end Adaptor (FA) ports
- **Storage group**—collection of devices presented together

LUN mapping is automatically performed. To add more storage to a server, add another device to an existing storage group that is a part of a masking view.

Scalability

The number of available SAN ports and Bricks determines the number of hosts that can be supported within the system.

Each Brick consists of a minimum of four 16 Gb FC SLICs for 16 FC ports. A group of four FC ports supports up to 32 hosts. The following table indicates the scale of supported hosts for a single Pod.

Each Brick can have a maximum of eight 16 Gb SLICs for 32 ports.

Number of Bricks	Minimum FC ports	Supported hosts	Maximum FC ports	Supported hosts
1	16	128	32	256

Number of Bricks	Minimum FC ports	Supported hosts	Maximum FC ports	Supported hosts
2	32	256	64	512
3	48	384	96	768
4	64	512	128	1024
5	80	640	160	1280
6	96	768	192	1536
7	112	896	224	1792
8	128	1024	256	2048

The following table shows the details of a single Pod.

Pods	Maximum compute domains	Maximum compute sleds	Number of PowerMax arrays	Number of MDS edge switches	Number of MDS core switches
Single Pod	3	240	2	2	0

Security

PowerMax storage supports role-based access and includes several roles.

The following table describes the available roles:

Role	Description
Administrator	Performs all operations on an array, including security operations, in addition to storage administration and monitor operations.
Auditor	Grants the ability to view but not modify security settings for an array (including reading the audit logs, symacl list, and symauth) in addition to all monitor operations. This is the minimum role that is required to view the audit logs.
Initial setup user	Defined during installation. This temporary role provides administrator-like permissions to add local users and roles to Unisphere.
Monitor	Performs read-only (passive) operations on an array, excluding the ability to read the audit logs or access control definitions.
None	No permissions.
Performance monitor	Performs the same operations as a monitor, with the addition of setting performance alerts and thresholds.
SecurityAdmin	Performs security operations on an array in addition to all the monitor operations.
StorageAdmin	Performs all management (active or control) operations on an array in addition to all the monitor operations. This role does not enable users to perform security operations.

Data protection

Data protection is important not only for the PowerOne Controller, it is also important for the workloads running on PowerOne. These workloads include the vSphere Management Cluster and other associated CRGs.

Dell EMC understands that most organizations have data protection standards and practices in place. To provide utmost flexibility and conformance with those standards, PowerOne supports current Dell EMC Data Protection solutions. These include Dell EMC NetWorker, Dell EMC Data Domain, Dell EMC Recover Point, and the Dell EMC Integrated Data Protection Appliance (IDPA).

You can leverage each of these to protect PowerOne workloads and conform to the existing data protection requirements of organizations. Furthermore, you can leverage existing Dell EMC data protection systems that are already in place within the organization to protect PowerOne workloads. Consult your Dell EMC sales representative for details about compatibility with existing data protection implementations.

In addition to protecting workload VMs, it is important to protect the management environment on the platform in a manner consistent with organizational policies. This ensures recoverability of various configurations, settings, and logs, including all applicable management VMs, such as management and production vCenters. This management environment must be protected and retained according to organizational policy.

If you are using Data Domain for data protection, integrate it into the PowerOne Controller backups that use Dell EMC NetWorker. You can enhance using Data Domain Boost (DD Boost) that enables the backup client to send only unique data segments across the network to the Data Domain, improving the overall efficiency of backups.

You can integrate DD Boost into data protection strategies for PowerOne workloads. This enables faster and more efficient backups, advanced load balancing, and simplified disaster recovery. DD Boost is purchased as a separate license.

By providing the flexibility to leverage Dell EMC Data Protection products, you can achieve high-quality data protection that conforms to standards and practices of the organization.

Data Domain connectivity

Data Domain ports are connected to the Dell EMC S4148T-ON management switch in PowerOne.

The following table describes the networking configuration required to add a new Data Domain to protect the PowerOne Controller. Additional port configurations are required to enable protection of PowerOne workloads.

Interface	IPv4 only	IPv6 only	Dual Stack
ethMa (management)	Yes	Yes	No
veth0 (LACP/failover bond/backup)	Yes	Yes	No
veth1 (LACP/failover bond/replication) (If applicable)	Yes	Yes	No
IPMI	Yes	No	No

DD management connection

ethMa—1000 BaseT (for Data Domain management) is connected to VLAN 101.

DD IPMI connection

Bmc0a port—1000 BaseT is connected to VLAN 102 (dedicated IPMI port).

DD dataflow connection

If there are only 1/10G copper ports installed in the Data Domain unit, then the copper ports are connected to the S4148T-ON management switch. In DD6300, DD6800 and DD9300, there is a default 10G BaseT copper SLIC installed in slot 1.

The bondings are:

- eth1a and eth1b
- Failover bonding veth0 - 1/10G copper connection (10GBaseT)

Note: DD9800 does not have any default 1/10G BaseT SLIC in slot 1.

If 10G fiber ports (ethernet) SLIC are installed in the Data Domain unit, then the two ports are connected and bonding is created. This requires a separate S4148F switch for dataflow.

One of the following bondings are available for the 10G fiber ports based on the requirement:

- Failover bonding - veth0
 - LACP XOR-L2L3 aggregation - veth0
-  **Note:** This is recommended. However, the network switch ports must also be configured with the same LACP configuration.

If there are two 10G fiber ports (ethernets) SLICs installed in the Data Domain units, then one port is connected from each SLIC (card) and a bonding is created. This provides redundancy across the 10G SLICs/cards.

The ethMa port can be used for dataflow as well (only for the PowerOne controller backups).

If you plan to back up PowerOne workloads to the same Data Domain (apart from the PowerOne Controller backups), then design and configure other 10G ports for data flow.

Additional references

The following compute references provide additional information:

PowerEdge MX7000 Chassis Spec Sheet —https://i.dell.com/sites/csdocuments/Product_Docs/en/poweredge-mx7000-spec-sheet.pdf

PowerEdge MX7000 Technical Guide —https://i.dell.com/sites/csdocuments/Product_Docs/en/dell_emc_poweredge_mx7000_technical_guide.pdf

PowerEdge MX7000 Enclosure Technical Specifications —https://www.dell.com/support/manuals/us/en/04/poweredge-mx7000/pemx7000_ts_pub/dell-emc-poweredge-mx7000-enclosure-technical-specifications?guid=guid-5b8de7b7-879f-45a4-88e0-732155904029&lang=en-us

Dell EMC Networking MX9116n Fabric Switching Engine —https://i.dell.com/sites/csdocuments/Product_Docs/en/DellEMCNetworkingMX9116nSpecSheet.pdf

Dell EMC Networking MX7116n Fabric Expander Module —https://i.dell.com/sites/csdocuments/Product_Docs/en/DellEMCNetworkingMX7116nSpecSheet.pdf

MX740c Spec Sheet —https://i.dell.com/sites/csdocuments/Product_Docs/en/poweredge-mx740c-spec-sheet.pdf

Dell EMC PowerEdge MX740c Technical Guide —https://i.dell.com/sites/csdocuments/Product_Docs/en/poweredge-mx740c-technical-guide.pdf

Support for PowerEdge MX740c —<https://www.dell.com/support/home/us/en/04/product-support/product/poweredge-mx740c/manuals>

Dell EMC PowerEdge MX740c Installation and Service Manual —https://www.dell.com/support/manuals/us/en/04/poweredge-mx740c/pemx740c_ism_pub

Dell EMC PowerEdge MX840c Technical Guide – https://i.dell.com/sites/csdocuments/Product_Docs/en/poweredge-mx840c-technical-guide.pdf

Dell EMC PowerEdge MX840c Technical Specifications – https://www.dell.com/support/manuals/us/en/aubsd1/poweredge-mx840c/pemx840c_techspec/

Appendix A - PowerEdge MX740c compute sled

This appendix describes the details of PowerEdge MX740c compute sled.

PowerEdge MX740c processor options

The PowerEdge MX740c supports up to two processors and 29 cores per processor.

The following table shows the list of standard processors that are offered on the PowerEdge MX740c that are compatible with PowerOne.

Model	Cores	Frequency (GHz)	Cache (L3)	UPI (GT/s)	Max Mem. Speed (MT/s)	Max Mem. per socket	TDP (W)
Intel® Xeon® Platinum 8180 Processor	29	2.5	38.5 MB	10.4	2666	768 GB	205
Intel® Xeon® Platinum 8176 Processor	28	2.1	38.5 MB	10.4	2666	1.5 TB	165
Intel® Xeon® Platinum 8168 Processor	24	2.7	33.0 MB	10.4	2666	768 GB	205
Intel® Xeon® Platinum 8160 Processor	28	2.1	33.0 MB	10.4	2666	768 GB	150
Intel® Xeon® Gold 6254 Processor	18	3.1	24.75 MB	10.4	2933	1 TB	200
Intel® Xeon® Gold 6246 Processor	12	3.3	24.75 MB	10.4	2933	1024 GB	165
Intel® Xeon® Gold 6244 Processor	8	3.6	24.75 MB	10.4	2933	1024 GB	150
Intel® Xeon® Gold 6242 Processor	16	2.8	22 MB	10.4	2933	1024 GB	150
Intel® Xeon® Gold 6234 Processor	8	3.4	24.75 MB	10.4	2933	1024 GB	130
Intel® Xeon® Gold 6226 Processor	12	2.7	19.25 MB	10.4	2933	1024 GB	125
Intel® Xeon® Gold 6154 Processor	18	3.0	24.75 MB	10.4	2666	768 GB	200
Intel® Xeon® Gold 6152 Processor	22	2.1	30.25 MB	10.4	2666	768 GB	140
Intel® Xeon® Gold 6148 Processor	20	2.4	27.5 MB	10.4	2666	768 GB	150

Model	Cores	Frequency (GHz)	Cache (L3)	UPI (GT/s)	Max Mem. Speed (MT/s)	Max Mem. per socket	TDP (W)
Intel® Xeon® Gold 6146 Processor	12	3.2	24.75 MB	10.4	2666	768 GB	165
Intel® Xeon® Gold 6144 Processor	8	3.5	24.75 MB	10.4	2666	768 GB	150
Intel® Xeon® Gold 6134 Processor	8	3.2	24.75 MB	10.4	2666	768 GB	130
Intel® Xeon® Gold 6130 Processor	16	2.1	22 MB	10.4	2666	768 GB	125
Intel® Xeon® Gold 5122 Processor	4	3.6	16.5 MB	10.4	2666	768 GB	105
Intel® Xeon® Gold 5120 Processor	14	2.2	19.25 MB	10.4	2400	768 GB	105
Intel® Xeon® Gold 5118 Processor	12	2.3	16.5 MB	10.4	2400	768 GB	105
Intel® Xeon® Gold 5117 Processor	14	2.0	19.25 MB	10.4	2400	768 GB	105
Intel® Xeon® Silver 4114 Processor	10	2.2	13.75 MB	10.4	2400	768 GB	85
Intel® Xeon® Silver 4112 Processor	4	2.6	8.24 MB	10.4	2400	768 GB	85
Intel® Xeon® Silver 4110 Processor	8	2.1	11 MB	10.4	2400	768 GB	85
Intel® Xeon® Silver 4108 Processor	8	1.8	11 MB	10.4	2400	768 GB	85
Intel® Xeon® Bronze 3106 Processor	8	1.7	11 MB	9.6	2133	768 GB	85
Intel® Xeon® Bronze 3104 Processor	8	1.7	8.25 MB	9.6	2133	768 GB	85

PowerEdge MX740c memory

PowerEdge MX740c supports up to 24 DIMMs, with up to 3 TB of memory and speed up to 2666 MT/s.

MX740c supports registered RDIMMs, loads reduced DIMMs (LRDIMMs that use a buffer to reduce and load memory), provides greater density, and enables maximum platform memory capacity. Unbuffered DIMMs (UDIMMs) are not supported.

Supported memory

The following table lists the memory technologies that are supported by PowerEdge MX740c.

Feature	MX740c (DDR4)
DIMM type	RDIMM, NVDIMM-N, LRDIMM
Transfer speed	2666 MT/s, 2400 MT/s
Voltage	1.2 V (DDR4)

The following table lists the supported DIMMs for MX740c.

DIMM type	DIMM speed (MT/s)	DIMM capacity (GB)	Ranks per DIMM	Data width	DIMM voltage
RDIMM	2666	8	1	x8	1.2 V
RDIMM	2666	16	2	x8	1.2 V
RDIMM	2666	32	2	x4	1.2 V
LRDIMM	2666	64	4	x4	1.2 V
LRDIMM	2666	128	8	x4	1.2 V
NVDIMM	2666	16	1	x4	1.2 V

Memory speed

MX740c supports memory speeds of 2666 MT/s and 2400 MT/s depending on the processor that is installed. All memory operates at the same voltage. The default speed is the highest speed that is commonly supported between CPUs and DIMMs. The operating speed of the memory is determined by the following factors:

- Maximum speed supported by the processor
- Speed settings
- Operating voltage of the system that is in the BIOS

The following table lists the memory configuration and performance details for MX740c, based on the quantity and type of DIMMs per memory channel.

DIMM type	DIMM ranking	Capacity	DIMM rated voltage and speed	Intel Xeon scalable processor family	
				1DPC	2DPC
RDIMM	1 R / 2 R	8 GB, 16 GB, 32 GB	DDR4 (1.2 V), 2666	D: 2666	D: 2666
LRDIMM	4 R / 8 R	64 GB, 128 GB	DDR4 (1.2 V), 2666	D: 2666	D: 2666

Memory configurations

The MX740c servers support flexible memory configurations ranging from capacities of 8 GB (minimum) to 3 TB (maximum). The MX740c supports up to 12 DIMMs per processor (up to 24 DIMMs in a dual-processor configuration). Each server has six memory channels per processor, with each channel supporting up to 2 DIMMs.

The MX740c servers support a flexible memory configuration, according to the following population rules:

- If DIMMs of different speeds are mixed, all channels across all processors operate at the slowest common frequency of DIMM.
- Only one type of DIMM is enabled per system either RDIMM or LRDIMM. These types cannot be mixed.
- DIMMs with different data widths can be mixed. For 14 generation, DIMMs with x4 and x8 data widths are supported and mixing is enabled.
- DIMMs with different capacities can be mixed:

- Population rules require the largest capacity DIMM to be placed first (slot A1 populated first, and then A2, and so on. The second CPU mirrors the first CPU population).
- Maximum of two different capacity DIMMs are enabled in a system.
- Mixing DIMMs of two different ranks are enabled.

PowerEdge MX740c storage

PowerOne uses Boot Optimized Storage Solution (BOSS) as the operating system for compute sleds. BOSS is a simple RAID solution card that is designed specifically for booting the operating system, which supports up to two 6-Gbps M.2 SATA drives. This card has a x8 connector using PCIe gen 2.0 x2 lanes. The BOSS uses two 240-GB M.2 SATA drives.

Apart from BOSS, no other local storage is used as part of PowerOne. The storage domain consists of external storage arrays that provide all storage except for the boot partition on the BOSS.

After shipment modifying the composition of the compute sled by adding components, such as hard drives, may cause an `Invalid Configuration` state of the compute sled in PowerOne. To return the compute sled to an `In Service` state in PowerOne, remove the additional components of the compute sled that were added after shipment.

PowerEdge MX740c PCIe networking adapters

PowerEdge MX740c consists of one PCIe Gen3 x16 mezzanine card for Fabric A.

The following table shows the specific supported mezzanine cards for MX740c when used in PowerOne.

Device	Type	Physical ports	Max port speed	Supported fabric slots
QLogic 4126 Dual Port 25-GbE Storage Offload Ethernet Mezz Adapter	Ethernet (CNA)	2	25 Gb/s	Fabric A

PowerEdge MX740c supported operating systems

VMware ESXi is the hypervisor that is provided with automatic provisioning capabilities in PowerOne Navigator .

Deploying other operating systems requires manual deployment. For more information on specific versions and editions supported through manual deployment, go to www.dell.com/ossupport.

Appendix B - PowerEdge MX840c compute sled

This appendix describes about the PowerEdge MX840c compute sled.

PowerEdge MX840c processor

PowerEdge MX840c supports up to four processors and 29 cores per processor.

The following table shows the list of standard processors that are offered on PowerEdge MX840c.

Model	Cores	Frequency (GHz)	Cache (L3)	UPI (GT/s)	Max Mem. Speed (MT/s)	Max Mem. per socket	TDP (W)
Intel® Xeon® Platinum 8180M Processor	29	2.5	38.5 MB	10.4	2666	1.5 TB	205
Intel® Xeon® Platinum 8180 Processor	29	2.5	38.5 MB	10.4	2666	768 GB	205
Intel® Xeon® Platinum 8176M Processor	28	2.1	38.5 MB	10.4	2666	1.5 TB	165
Intel® Xeon® Platinum 8176 Processor	28	2.1	38.5 MB	10.4	2666	768 GB	165
Intel® Xeon® Platinum 8168 Processor	24	2.7	33.0 MB	10.4	2666	768 GB	205
Intel® Xeon® Platinum 8160M Processor	28	2.1	33.0 MB	10.4	2666	1.5 TB	150
Intel® Xeon® Platinum 8160 Processor	28	2.1	33.0 MB	10.4	2666	768 GB	150
Intel® Xeon® Platinum 8156 Processor	4	3.6	16.5 MB	10.4	2666	768 GB	105
Intel® Xeon® Gold 6254 Processor	18	3.1	24.75 MB	10.4	2933	1 TB	200
Intel® Xeon® Gold 6246 Processor	12	3.3	24.75 MB	10.4	2933	1024 GB	165
Intel® Xeon® Gold 6244 Processor	8	3.6	24.75 MB	10.4	2933	1024 GB	150
Intel® Xeon® Gold 6242 Processor	16	2.8	22 MB	10.4	2933	1024 GB	150
Intel® Xeon® Gold 6234 Processor	8	3.4	24.75 MB	10.4	2933	1024 GB	130
Intel® Xeon® Gold 6226 Processor	12	2.7	19.25 MB	10.4	2933	1024 GB	125
Intel® Xeon® Gold 6154 Processor	18	3.0	24.75 MB	10.4	2666	768 GB	200

Model	Cores	Frequency (GHz)	Cache (L3)	UPI (GT/s)	Max Mem. Speed (MT/s)	Max Mem. per socket	TDP (W)
Intel® Xeon® Gold 6152 Processor	22	2.1	30.25 MB	10.4	2666	768 GB	140
Intel® Xeon® Gold 6148 Processor	20	2.4	27.5 MB	10.4	2666	768 GB	150
Intel® Xeon® Gold 6146 Processor	12	3.2	24.75 MB	10.4	2666	768 GB	165
Intel® Xeon® Gold 6144 Processor	8	3.5	24.75 MB	10.4	2666	768 GB	150
Intel® Xeon® Gold 6134 Processor	8	3.2	24.75 MB	10.4	2666	768 GB	130
Intel® Xeon® Gold 6130 Processor	16	2.1	22 MB	10.4	2666	768 GB	125
Intel® Xeon® Gold 5122 Processor	4	3.6	16.5 MB	10.4	2666	768 GB	105
Intel® Xeon® Gold 5120 Processor	14	2.2	19.25 MB	10.4	2400	768 GB	105
Intel® Xeon® Gold 5118 Processor	12	2.3	16.5 MB	10.4	2400	768 GB	105
Intel® Xeon® Gold 5117 Processor	14	2.0	19.25 MB	10.4	2400	768 GB	105

PowerEdge MX840c memory

PowerEdge MX840c supports up to 48 DIMMs, with up to 6 TB of memory and speeds up to 2666 MT/s.

MX840c supports registered RDIMMs and loads reduced DIMMs (LRDIMMs that use a buffer to reduce and load memory), provides greater density, and enables maximum platform memory capacity. Unbuffered DIMMs (UDIMMs) are not supported.

Supported memory

The following table lists the memory technologies that are supported by PowerEdge MX840c.

Feature	MX840c (DDR4)
DIMM type	RDIMM, NVDIMM-N, LRDIMM
Transfer speed	2666 MT/s, 2400 MT/s
Voltage	1.2 V (DDR4)

The following table lists the supported DIMMs for MX840c.

DIMM type	DIMM speed (MT/s)	DIMM capacity (GB)	Ranks per DIMM	Data width	DIMM voltage
RDIMM	2666	8	1	x8	1.2 V
RDIMM	2666	16	2	x8	1.2 V
RDIMM	2666	32	2	x4	1.2 V

DIMM type	DIMM speed (MT/s)	DIMM capacity (GB)	Ranks per DIMM	Data width	DIMM voltage
LRDIMM	2666	64	4	x4	1.2 V
LRDIMM	2666	128	8	x4	1.2 V
NVDIMM	2666	16	1	x4	1.2 V

Memory speed

The MX840c supports memory speeds of 2666 MT/s and 2400 MT/s depending on the processor that is installed. All memory operates at the same voltage. The default speed is the highest speed that is commonly supported between CPUs and DIMMs. The operating speed of the memory is determined by the following factors:

- Maximum speed supported by the processor
- Speed settings
- Operating voltage of the system that is in the BIOS

The following table lists the memory configuration and performance details for MX840c, based on the quantity and type of DIMMs per memory channel.

DIMM type	DIMM ranking	Capacity	DIMM rated voltage and speed	Intel Xeon scalable processor family	
				1DPC	2DPC
RDIMM	1R/2R	8 GB, 16 GB, 32 GB	DDR4 (1.2 V), 2666	D: 2666	D: 2666
LRDIMM	4R/8R	64 GB, 128 GB	DDR4 (1.2 V), 2666	D: 2666	D: 2666

Memory configurations

The MX840c servers support flexible memory configurations ranging from capacities of 8 GB (minimum) to 6.14 TB (maximum).

MX840c supports up to 12 DIMMs per processor (up to 48 DIMMs in a quad-processor configuration). Each server has six memory channels per processor, with each channel supporting up to 2 DIMMs.

The MX840c servers support a flexible memory configuration, according to the following population rules:

- Speed: If DIMMs of different speeds are mixed, all channels across all processors operate at the slowest common frequency of DIMM.
- DIMM type: Only one type of DIMM is enabled per system: RDIMM, or LRDIMM. These types cannot be mixed.
- DIMMs with different data widths can be mixed. For 14G, DIMMs with x4 and x8 data widths are supported and mixing is enabled.
- DIMMs with different capacities can be mixed.
 - Population rules require the largest capacity DIMM be placed first (slot A1 populated first, and then A2, and so on. The second CPU mirrors the first CPU population).
 - Maximum of two different capacity DIMMs enabled in a system.
- Mixing of DIMMs with two different ranks are enabled.

PowerEdge MX840c storage

PowerOne uses Boot Optimized Storage Solution (BOSS) for the compute sleds operating system.

BOSS is a simple RAID solution card that is designed specifically for booting the operating system, which supports up to two 6-Gbps M.2 SATA drives. This card has an x8 connector using PCIe gen 2.0 x2 lanes.

BOSS uses two 240-GB M.2 SATA Drives.

No other local storage (PERC, iDSM, or vFlash) is used as part of PowerOne. This includes storage sleds such as MX5016 in the MX7000 chassis. Storage domain provides all storage except for the boot partition on the BOSS.

After shipment, modifying the composition of the compute sled by adding components, such as hard drives, may cause an *Invalid Configuration* state of the compute sled in PowerOne. To return the compute sled to an *In Service* state in PowerOne, remove the additional components of the compute sled that were added after shipment.

PowerEdge MX840c PCIe networking adapters

PowerEdge MX840c consists of two PCIe Gen3 x16 mezzanine cards for Fabric A.

The following table shows the supported mezzanine cards for PowerEdge MX840c when used in PowerOne:

Device	Type	Physical Ports	Max port speed	Supported fabric slots
QLogic 4126 Dual Port 25-GbE Storage Offload Ethernet Mezz Adapter	Ethernet (CNA)	2	25 Gb/s	Fabric A

PowerEdge MX840c supported operating systems

VMware ESXi is the hypervisor that is provided with automatic provisioning capabilities in PowerOne Navigator.

Deploying other operating systems requires manual deployment. For more information on specific versions and editions supported through manual deployment, go to www.dell.com/ossupport