

Dell EMC PowerEdge MX7000

Technical Guide

Notes, cautions, and warnings

 **NOTE:** A NOTE indicates important information that helps you make better use of your product.

 **CAUTION:** A CAUTION indicates either potential damage to hardware or loss of data and tells you how to avoid the problem.

 **WARNING:** A WARNING indicates a potential for property damage, personal injury, or death.

Chapter 1: System overview	6
Introduction.....	6
Featured technologies.....	6
Chapter 2: System features.....	8
Product comparison.....	8
Specifications.....	9
Chapter 3: Chassis view and features.....	10
Front view of the enclosure.....	10
Left control panel.....	11
Right control panel.....	13
Back view of the enclosure.....	13
Chapter 4: Internal Architecture.....	15
Infrastructure boards.....	15
System interconnection.....	16
Chapter 5: Input output fabrics.....	17
Input output fabrics.....	17
Chapter 6: Storage fabrics.....	19
Storage fabric.....	19
Chapter 7: System Modules.....	20
Management Module (MM).....	20
Compute Sled.....	20
Storage Sled.....	20
I/O Modules.....	21
Power Supplies	21
Fans	21
Chapter 8: Rail kit.....	22
Chapter 9: Power, thermal, and acoustics	23
Power.....	23
Thermal.....	24
Acoustics.....	25
Chapter 10: OpenManage systems management.....	27
iDRAC with Lifecycle controller.....	27
Agent-free management.....	30
Agent-based management.....	30

OpenManage console.....	30
OpenManage systems management tools, utilities and protocols.....	31
OpenManage Integrations.....	31
OpenManage Connections.....	31
Chapter 11: Appendix A. Additional specifications.....	32
Population rules.....	32
PSU redundancy and population rules.....	33
Chassis dimensions.....	34
Chassis weight.....	34
Fan specifications.....	35
PSU specifications.....	35
Ports and connectors specifications.....	35
USB ports.....	35
Mini DisplayPort.....	35
PowerEdge MX modules ports and connectors.....	36
PowerEdge MX740c.....	36
PowerEdge MX840c.....	36
MX7116n Fabric Expander Module.....	36
MX9116n Fabric Switching Engine.....	36
MX9002m Management Module.....	37
MX5108n Ethernet Switch.....	37
MXG610s Fibre Channel Switch.....	37
PowerEdge MX 10GBASE-T Ethernet Pass-Through Module.....	37
PowerEdge MX 25 Gb Ethernet Pass-Through Module.....	37
Video specifications.....	38
Environmental specifications.....	38
Standard operating temperature.....	39
Expanded operating temperature.....	39
Particulate and gaseous contamination specifications	40
Chapter 12: Appendix B. Standards compliance.....	42
Chapter 13: Appendix C. Additional resources.....	43
Chapter 14: Appendix D. Support and deployment services.....	44
Dell EMC ProDeploy Enterprise Suite	44
Dell EMC ProDeploy Plus.....	44
Dell EMC ProDeploy.....	45
Dell EMC Basic Deployment.....	45
Dell EMC Residency Services.....	45
Deployment services.....	45
Dell EMC Remote Consulting Services.....	45
Dell EMC Data Migration Service.....	45
ProSupport Enterprise Suite.....	45
ProSupport Plus.....	46
ProSupport.....	46
ProSupport One for Data Center.....	47
Support Technologies.....	47

Additional professional services.....	48
Dell Education Services.....	48
Dell EMC Global Infrastructure Consulting Services.....	48
Dell EMC Managed Services.....	48

System overview

The Dell EMC PowerEdge MX portfolio delivers a fully managed, high performance system that will free up valuable IT resources and personnel so you can focus on innovation. Break free from the bounds of technology silos and routine, daily and time consuming operational management to realize your IT and digital business transformations.

PowerEdge MX, a unified, high performance kinetic infrastructure, provides the agility, resiliency and efficiency to optimize a wide variety of traditional and new, emerging data center workloads and applications. With its kinetic architecture and agile management the MX portfolio dynamically configures compute, storage and fabric, increases team effectiveness and accelerates operations. Its responsive design delivers the innovation and longevity customers of all sizes need for their IT and digital business transformations. PowerEdge MX ecosystem consists of a new chassis infrastructure, compute sleds, fabric switches, and a storage sled, all managed by Dell EMC OpenManage Enterprise-Modular Edition.

Topics:

- [Introduction](#)
- [Featured technologies](#)

Introduction

The MX7000 and MX-series portfolio delivers a fully managed, high performance system that will free up valuable IT resources and personnel so you can focus on innovation. Break free from the bounds of technology silos and routine, daily and time consuming operational management to realize your IT and digital business transformations. The new PowerEdge MX, a unified, high performance kinetic infrastructure, provides the agility, resiliency and efficiency to optimize a wide variety of traditional and new, emerging data center workloads and appliances. With its kinetic architecture and agile management the MX portfolio dynamically configures compute, storage and fabric, increases team effectiveness and accelerates operations. Its responsive design delivers the innovation and longevity customers of all sizes need for their IT and digital business transformations.

The PowerEdge MX infrastructure provides:

Flexible architecture

- A flexible architecture – Non-disruptive provisioning, on-demand allocation of compute, storage and networking resource pools.
- A scalable fabric – Cost-effective multi-chassis architecture with a broad array of open networking options and upgrade simplicity for future I/O flexibility.
- Granular storage – Dense, highly flexible, hot swappable, scale-out direct attached storage sled with easy front access.

Agile management

- End-to-end lifecycle management and single point authentication for all devices from a single interface.
- Simplified set-up/updates with no specialized training needed, and with multiple at-the-box management options.
- An operational template methodology and comprehensive Rest API.

Responsive design

- An industry leading fabric, system thermal architecture, mechanical design and control algorithms for dense configurations with future compatibility.
- A hardened design to protect, detect and recover underlying infrastructure from cyber attacks.

Featured technologies

The following are the new technologies that are featured on the PowerEdge MX7000:

Table 1. Featured technologies

New technology	Detailed description
Form factor	The MX7000 has all new chassis, sled, and I/O module form factors

Table 1. Featured technologies (continued)

New technology	Detailed description
I/O Fabrics	The MX7000 support flexible I/O Fabrics
Direct Orthogonal connection	Direct Orthogonal Compute to I/O connections allowing for future technology upgrades
System management solution	New Systems Management solution including x86 management processor option
KVM module	KVM solution integrated into Management processor

System features

The PowerEdge MX740c compute sled is a unique design for the PowerEdge MX7000 chassis and is not compatible with other PowerEdge modular chassis.

The MX7000 chassis has the following features:

- Sleds
 - 8 single-width or 4 double-width sleds
 - Double wide sleds in slots 1-2, 3-4, 5-6, 7-8
- Power supply units
 - Up to 6 front loading power supply units
- Fans
 - 4 x 60 mm front fan modules
 - 5 x 80 mm rear fan modules
- Right control panel ear
 - Power button and LED indicators
 - 2 USB Type-A and 1 USB type Micro-AB
 - 1 Mini DisplayPort
- Left control panel ear - three configurations
 - LCD with QuickSync
 - LCD without QuickSync
 - LED indicators only

For additional information, please refer to the PowerEdge MX7000 Technical Guide.

Topics:

- [Product comparison](#)
- [Specifications](#)

Product comparison

Table 2. Feature comparison

Feature	MX7000	M1000e
Rack height	7U rack chassis enclosure	10U blade chassis enclosure
Sled orientation	Vertical	Vertical
Sled support	Up to: Eight single-width sleds or Four double-width sleds Up to seven Storage sleds can be populated in the enclosure.	Up to: 16 half-height blades or 8 full-height blades or 32 quarter-height blades
Input output modules	3+3 (A/B/C) A/B direct-IO interconnect C is Midplane-based interconnect	3+3 (A/B/C) Midplane-based interconnect A reserved for Ethernet
Power Supplies	6 x 3 kW	6 x 3 kW
KVM	Integrated with the Management module	Discrete design

Specifications

Table 3. Technical specifications

Features	Specifications
Form factor	7U rack
Sled orientation	Vertical
Sled support	Up to 8 standard-height, or 4 double-wide, or 16 half-height sleds
I/O modules	3+3 (A/B/C). A and B direct I/O Ethernet interconnect. C is midplane-based fiber channel or SAS interconnect
Power supplies	6 x 3 KW
KVM	Integrated on the Management Module

Chassis view and features

The MX740c is a full-height, single-width, two socket compute sled for the MX7000 chassis.

Topics:

- Front view of the enclosure
- Back view of the enclosure

Front view of the enclosure




Figure 1. Front view of the enclosure

- | | |
|------------------------------|------------------------------|
| 1. Left control panel | 2. Single-width compute sled |
| 3. Sled blank | 4. Front fan (4) |
| 5. Double-width compute sled | 6. Single-width storage sled |
| 7. Right control panel | 8. Information tag |
| 9. Power supply unit (6) | |

Left control panel



Figure 2. Left control panel - Status LED

Table 4. Left control panel - LED indicator description

Indicator	Description	Status	
1	System health	Blinks amber for 2 seconds and is OFF for 1 second when the chassis health has degraded. By default, the LED is unlit.	
2	System temperature	Blinks amber for 2 seconds and is OFF for 1 second when a thermal fault exists on the enclosure. By default, the LED is unlit. NOTE: A thermal fault includes excessive ambient temp, I/O modules thermal status, PSU thermal status, and fan status.	
3	I/O module health	Blinks amber for 2 seconds and is OFF for 1 second when an I/O module is faulty. By default the LED is unlit.	
4	Fan health	Blinks amber for 2 seconds and is OFF for 1 second when a front or rear mounted fan fails or has a warning. By default, the LED is unlit.	
5	Stack or group	Indicates that the enclosure is a member of a group.	
6	LED status bar	Indicator status	Description
		Solid blue	Indicates that the enclosure is healthy.
		Blinking blue	Indicates that the system ID mode is active.
		Blinking amber	Indicates that the system is experiencing a fault.
7	System ID button	Allows you to identify the system or the installed sleds.	




Figure 3. Left control panel - LCD options

Table 5. Left control panel - LCD panel description

Indicator	Description	Status						
1	LCD with Quick Sync	LCD enabled with Quick Sync module						
2	LCD without Quick Sync	LCD without Quick Sync module						
3	System ID indicator on LCD panel	This option is a button/indicator on the LCD panel to identify the chassis, or choose specific sleds to identify.						
4	Settings	This option button provides access to the inventory and configuration data of the MX7000 enclosure. It includes the Network Settings, System Information, (Model, Asset Tag, Service Tag), and Language Settings.						
5	Optional QuickSync indicator (Only for LCD with QuickSync 2.0)	<p>Enables access to QuickSync related controls and connection information.</p> <p>NOTE: QuickSync feature allows you to manage your system using mobile devices. This feature is only available on certain configurations.</p> <p>NOTE: If not ordered at the time of purchase, the QuickSync module will not be available on the enclosure.</p>						
6	System alerts indicator	<table border="1"> <thead> <tr> <th>System ID Indicator status</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>Solid green</td> <td>The chassis has no degraded or critical alerts.</td> </tr> <tr> <td>Solid amber</td> <td>The chassis has critical or degraded health alerts.</td> </tr> </tbody> </table>	System ID Indicator status	Description	Solid green	The chassis has no degraded or critical alerts.	Solid amber	The chassis has critical or degraded health alerts.
		System ID Indicator status	Description					
		Solid green	The chassis has no degraded or critical alerts.					
		Solid amber	The chassis has critical or degraded health alerts.					
<p>NOTE: This option button/indicator shows an amber colored alert icon and a combined critical and degraded alert count. Pressing the button takes the user to the alert details menu.</p>								
7	LCD activation button/ System ID indicator/ Identification indicator	<p>Allows you to identify the enclosure.</p> <p>NOTE: Press the button to activate the LCD.</p>						
		<table border="1"> <thead> <tr> <th>System ID Indicator status</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>Blinking blue</td> <td>System ID is active.</td> </tr> <tr> <td>Blinking amber</td> <td>Chassis alerts are present.</td> </tr> </tbody> </table>	System ID Indicator status	Description	Blinking blue	System ID is active.	Blinking amber	Chassis alerts are present.
		System ID Indicator status	Description					
		Blinking blue	System ID is active.					
Blinking amber	Chassis alerts are present.							
8	Error indicator	The error indicator is displayed on the LCD when there are any critical/warning alerts on the enclosure.						
9	Optional Quick Sync wireless status indicator	Displays the connection status of the enclosure with any QuickSync enabled device.						

Right control panel



Figure 4. Right control panel

1. Power button
2. USB 2.0 port (2)
3. Mini display port

Back view of the enclosure




Figure 5. Back view of the enclosure

- | | |
|-------------------------------------|-----------------------------------|
| 1. Slot for Fabric A1 | 2. Slot for Fabric A2 |
| 3. Rear fans (5) | 4. Slot for Fabric B1 |
| 5. Slot for Fabric B2 | 6. Slot for Fabric C2 |
| 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

Internal Architecture

The following sections contain information about the internal architecture of Dell EMC PowerEdge MX7000 system.

The following diagram is a simplified view of the MX7000 system, showing some of the modules and their interconnection.

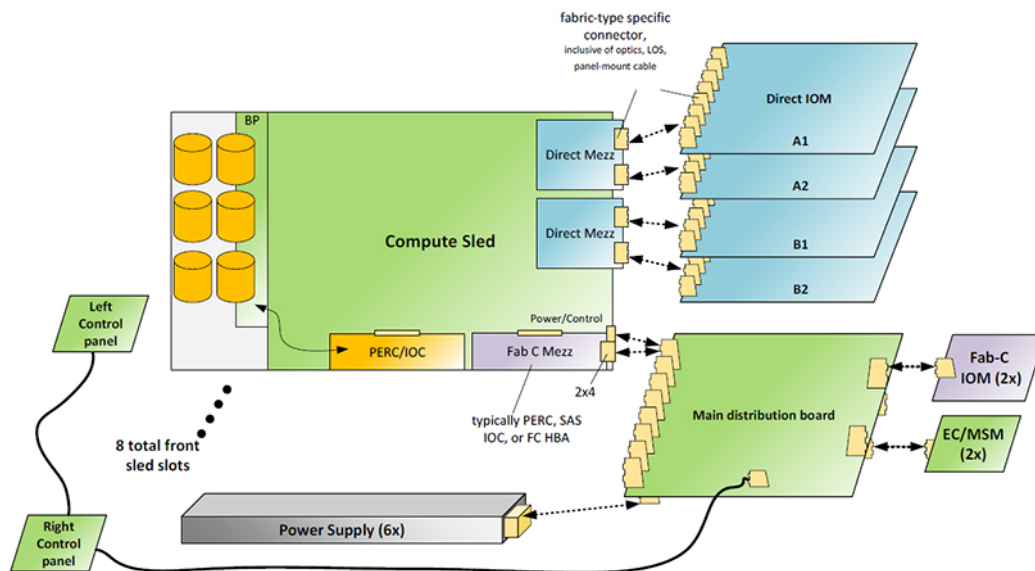


Figure 6. Internal Architecture

Topics:

- [Infrastructure boards](#)
- [System interconnection](#)

Infrastructure boards

Main distribution board

The main distribution board is the most complex of all the internal infrastructure boards. There are no active components, however there are many signals interconnected on this board. The main distribution board is horizontally placed near the chassis floor. From the front of the chassis, the six power supplies plug into connectors on the bottom of the main distribution board. Also from the front of the chassis the eight compute sleds plug into this board for power, management signals, and other fabric connections. From the rear of the chassis the Management Modules (MMs) plug in to this board to provide overall chassis management and control. Also from the rear of the chassis the two fabric I/O Modules (Fabric - C IOMs) plug into this board.

Vertical power distribution board

The vertical power distribution board is above the main distribution board, oriented perpendicular to the chassis floor. There are no active components, but many signal connections on this board. This board connects to the main distribution board with a flex cable, and power and ground bus bars. From the front of the chassis, the four cooling fans plug into the vertical power distribution board. Also from the rear of the chassis the four fabric I/O Modules (Fab-A/B IOMs) plug into this board for power, management signals, and other fabric connections. A cable connects this board to the rear fan board.

Rear fan board

The rear fan board, or 80 mm Fan Interposer, sits parallel to the chassis floor. The rear fans plug into this board from the rear of the chassis. This board connects internally to the vertical power distribution board, through a signal and power cable.

System interconnection

Power supply AC connection

The PSUs are accessed from the front, however the AC power cable connection is at the rear of the chassis. There are six C22 inlet connectors, three on the rear left, and three on the rear right side. Each PSU has an associated AC good LED located next to its associated plug. If a PSU is not present, plugging in an AC cord has no effect. If a PSU is present, and the AC power source is providing the appropriate power signal, the PSU drives the LED ON.

The C22 Inlet connectors allow for higher exhaust temperatures. The connectors are different than previous PowerEdge offerings, a C20 to C21 jumper cord is required to connect to the inlet connectors.

Internal fabrics

This infrastructure has numerous communication paths or fabrics interconnecting the modules. The management fabric, referred to as Fabric-D 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, an improvement over previous infrastructures which have run at 100 Mb/sec. A 10Gbit/sec link spans the Main DB between the two MM slots for communication between those modules.

A second fabric in the infrastructure is Fabric-C. This fabric connects between the IOM C1/C2 modules and each front sled slot. This fabric is intended for storage applications, for example SAS. Each front slot Fabric-C connection has two links, each x4 in width to each of the IOMs, IOM_C1 and IOM_C2. Today, this fabric supports 12 Gb/sec SAS and Fibre Channel at 32 Gb/sec.

USB

Each MM has a USB 2.0 link to the control panel. Only one port is operational at a time, based on which MM is active.

Front control panel and power control interfaces

Each MM displays the video to the front control panel Mini DisplayPort. The active MM drives the video. In addition to the DP video, there are also individual links for USB from each MM to the right control panel.

Input output fabrics

The following sections contain information about Dell EMC PowerEdge MX7000 Input output fabrics.

Topics:

- [Input output fabrics](#)

Input output fabrics

The internal connection of the main I/O subsystem in MX7000 enclosure is referred to as fabric A and B. These fabrics are connected using a direct orthogonal connection between sleds in the front of the chassis and I/O fabric modules in the rear. The mating connection on the sleds is on the mezzanine cards, which can install into location A or B on the sled.

Each fabric A and B allows for two IOMs (that is, A1 and A2, or B1 and B2) for redundancy. Each IOM connects to all 8 compute nodes through the orthogonal connectors. The connection from the IOM to each compute node (The mezzanine card on the compute node) is a 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.

Standalone pass through modules

Two stand alone Ethernet pass through modules are supported in the initial release, a 10G (Base-T) and 25 GbE pass through module. The pass-through IOMs plug into the fabric A or B IOM slots in the chassis. These IOMs redrive, or pass through the link signals without any switching. Each shall provide 16 internal server-facing ports and 16 external user-facing ports.

i **NOTE:** Only the ports 1-8 are active, when compute sleds with dual-port mezzanine adapters are configured with the pass-through modules. The ports 9-16 require a quad-port mezzanine adapter.

Front sled to rear port mappings

The external port assignments of pass-thru modules are fixed due to the direct routing connection internal to the system and the IOM itself.

i **NOTE:** If storage or other peripheral sleds are installed that do not support fabric A/B connections, the corresponding ports on the I/O modules mapping to these slots are not active. This is also the case for any compute sled that does not have a mezzanine card that is installed for the particular IOM (A or B) in the system.

A standard size compute sled is configured with dual port fabric A/B mezzanine maps to a single port on each (A or B) I/O module.

Table 6. Compute sleds dual-port mapping

Compute sled	IOM x1 (top)	IOM x2 (bottom)
1	1	1
2	2	2
3	3	3
4	4	4
5	5	5
6	6	6
7	7	7
8	8	8

Table 7. Compute sleds with quad-port mezzanine cards mapping

Compute sled	IOM x1 (top)	IOM x2 (bottom)
1	1, 9	1, 9
2	2, 10	2, 10
3	3, 11	3, 11
4	4, 12	4, 12
5	5, 13	5, 13
6	6, 14	6, 14
7	7, 15	7, 15
8	8, 16	8, 16

Other I/O modules

In addition to the standalone pass through modules previously described, MX7000 enclosure supports the following IOMs:

- MX9116n
- MX7116n
- MX5108n
- MXG610s

Storage fabrics

The following sections contain information about Dell EMC PowerEdge MX7000 storage fabrics.

Topics:

- [Storage fabric](#)

Storage fabric

The internal connection of the storage subsystem in the MX7000 enclosure is referred to as Fabric C. This fabric is connected on the Main DB between sleds in the front of the MX7000 enclosure and I/O modules in the rear. The primary use cases for this fabric will be for SAS or Fibre Channel storage connectivity, however other technologies could utilize this interconnect in the future.

Each compute sled can support a Fabric-C Mezz card, which connects to each IOM through a x4 link. The SAS Fabric-C IOMs provide flexible SAS switching between compute sleds and internal storage sleds (MX5016s) connected to the MX7000 enclosure. The fiber channel IOMs provide connectivity between compute sleds and external storage.

Fabric-C IOM interconnect

There are several connections on the Main DB to enable communication between the IOMs. Between each pair of IOMs (C1 and C2), there is a link for intermodule communication. This link supports a x1 connection with each lane operating up to 10 Gbps per direction. Also, there are handshake signals between the two modules to facilitate redundancy/failover or other module to module communication. The usage of these interconnects depends on the IOM design and its particular requirements.

System Modules

The following sections contain information about Dell EMC PowerEdge MX7000 system modules.

Topics:

- [Management Module \(MM\)](#)
- [Compute Sled](#)
- [Storage Sled](#)
- [I/O Modules](#)
- [Power Supplies](#)
- [Fans](#)

Management Module (MM)

The MX9002m Management Module (MM), or MM, essentially controls the overall chassis power, cooling, and physical user interfaces such as the front panel. It also contains the Ethernet switch used by the management network (Fab-D) and external Ethernet ports. MX7000 supports two MM for redundancy. At least one MM is required to power up the system.

External Network RJ-45 Ports

The MM internal network switch supports dual 1000BASE-T uplink ports to the chassis management network. The management network provides Gigabit Ethernet (GbE) connectivity to a Sleds, and IOMs. Either uplink port may be used as a stacking port if multiple MM are daisy-chained to networking switches.

External Micro USB Serial Port

The MX9002m Management Module (MM) provides a Micro-USB port for serial access by an external host. The UART settings are 115K baud, 8 data bits, 1 stop bit.

System ID Button

The System ID button on the MM rear panel allows for identifying the system in a rack. When pressed, the front panel of the chassis will also blink in the identify pattern. This button also serves as the Health status LED indicator, displaying blue for health or blinking amber if the MM has a fault.

The rear mounted ID button ONLY illuminates the Chassis/MM ID button, it does not step through the other devices like the Left Control Panel ID Button.

Compute Sled

The MX7000 supports up to eight standard size or up to four double wide compute sleds. The MX7000 supports 14G Next-Generation Modular (NGM) compute sleds but does not support 13G, FX2 or M1000e form-factor sleds.

Storage Sled

The MX7000 can accept a storage sled in the front slot, using the same form factor of a compute sled. The MX5016s is an example of this type of sled. The storage sled shall use the Fabric C interconnect within the chassis to communicate with the IOMs.

NOTE: There is no preassigned mapping of a storage sled to a compute sled in the MX7000/MX5016s configuration. User may map the entire storage sled to any compute sled, or map specific hard drives within the storage sleds to different compute sleds.

I/O Modules

The I/O Modules (IOMs) support either a Fabric A/B interconnect or a Fab-C interconnect. The Fab A/B IOMs 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 PDB and the sleds. .

The Fab C IOMs span one half of the chassis width, installing side by side in the back of the chassis. These IOMs plug directly into the Main DB using three high density right-angle connectors.

Power Supplies

The MX7000 Power Supplies (PSUs) utilize a new front-loading form factor. Up to six PSUs plug into the front of the chassis. These mate directly with connectors on the Main Distribution Board (DB). The AC connections come from the rear of the chassis cabling to connections on the Power Distribution Board (PDB).

Fans

There are two different types of fan module used in the MX7000 chassis, 60mm fans in the front and 80mm fans in the rear. Each fan is enclosed in a mechanical assembly to make up the module. The fan modules are hot-swappable, with 3+1 redundancy for the front fans, and 4+1 redundancy for the rear fans. The system should always be populated with the full set of fans to support the airflow requirements of the chassis.

Rail kit

A static rail kit will be supported for the chassis. Due to the overall weight, there is no support for 2-post racks. A strain relief bar will be provided for power cables. This bar adds 95mm of additional depth to the system, and therefore works best in 1200mm depth racks.

Power, thermal, and acoustics

Topics:

- Power
- Thermal
- Acoustics

Power

Power Supply

The MX7000 Power Supplies will utilize a new front-loading form factor developed for the MX portfolio. Up to six PSUs plug into the front of the chassis. These mate directly with connectors on the Main DB. The AC connections come from the rear of the chassis cabling to connections on the PDB. The Platinum PSUs support 3000W output with high line AC input. For low line input the PSUs will output approximately 1400W. Low line will not be supported at RTS. If low line is ever supported, mixed high line and low line AC inputs in the same chassis will not be supported. The supplies are set nominally to output 12.2V DC at the output.



Figure 7. PSU indicators

1. PSU health indicator
2. AC supply status indicator
3. DC output status indicator

Table 8. PSU health indicator codes

PSU health indicator	Indicator state
PSU functioning normally	Green
PSU faulty	Blinking amber
PSU mismatch	ON for 1 second, and then 5 blinks and OFF (non-repeating cycle).

Table 9. AC indicator codes

AC indicator	Indicator state
AC source available	ON
AC source unavailable or power cable unplugged	OFF

Table 10. DC indicator codes

DC indicator	Indicator state
DC output available	ON
DC output unavailable	OFF

Thermal

The chassis is cooled from front to back, as with typical rack systems. There are separate paths for the exhaust, keeping the PSU exit air isolated from the sled exit air. The sled air flow is controlled by the rear 80 mm fans, while the PSU air flow is controlled by the PSU fans located toward the front of the PSUs.

Ambient Temperature Requirements

In general, the EMC PowerEdge MX7000 shall support an ambient operating temperature of 10° C to 35° C. Some module configurations (for example, compute sleds with special high wattage CPU SKUs) are not able to support 35° C ambient with full performance. And there are some modules and features that cannot be supported at Fresh Air ranges.

Table 11. Ambient Temperature restrictions

	30°C	35°C	40°C ~ 45°C Expanded Operating Temperature
Chassis (MX7000) Includes Fans, MM, and PSUs	No restriction	No restriction	No restriction
Fabric A/B Modules	No restriction	No restriction	Not Supported List: MX9116n
Fabric C Modules	No restriction	No restriction	No restriction
2S Compute (MX740c)	No restriction	200 W & 205 W CPUs experience performance reduction when configured with 6 hard drives	Not Supported Features: 150 W CPUs and above Other specific CPU SKU restrictions: 6146 (165 W / 12c) 6144 (150 W / 8c) 6134 (130 W / 8c) 6128 (115 W / 6c) 5122 (105 W / 4c) NVDIMM NVMe hard drive Mezz cards >30 W
4S Compute (MX840c)	No restriction	No restriction	Not Supported Features: 165 W CPUs and above NVDIMM NVMe hard drive Mezz cards >30 W
Storage Sled (MX5016s)	No restriction	No restriction	No restriction

Acoustics

Dell EMC PowerEdge MX7000 acoustical design is guided by the following principles:

- Deliver sound levels that are appropriate for end users' deployments
- Deliver sound quality (for example tones, hums) that does not distract those working in proximity to the product
- Tune fan control to quickly attain and maintain target fan speeds while minimizing oscillation
- Deliver efficient thermal design that meets reliability and performance requirements at the lowest fan speeds
- Minimize vibration for sound quality (for example no rattles) and reduce disturbances to rotational hard drive media

Deployment Environment

MX7000 infrastructure is designed for use in unattended data center environments; its acoustical output will be too loud for office areas in all configurations. Dell EMC provides options for the customer to reduce or limit acoustical output for attended data center applications, including: configuration restrictions, sound cap profile (iDRAC), and a hardware acoustical baffle.

Acoustical Performance

The following table shows acoustical performances of minimum, typical, and feature-rich configurations.

Table 12. Acoustical Performance configurations

Configuration ^a	Minimum	Typical-1	Feature Rich
Compute Sled	4x MX740c (85 W CPU)	4x MX740c (140 W CPU) 2x MX840c (165 W CPU)	6x MX740c (165 W CPU)
Storage Sled	None	None	2x MX 5016 s
FAB A/B	10 GbE PTM	LE IOM	MX9116n & MX7116n
FAB C	None	MXG610s	MX5000s

Table 13. Acoustical Performance

Acoustical Performance: Idle/ Operating @ 25°C Ambient				
LwA-UL ^b (Bels)	Idle ^c	6.4	7.1	7.3
	Storage Stressed	6.4	7.1	7.4
	CPU 50% Utilized	6.4	7.1	7.4
Prominent tones ^d		None	None	None

NOTE:

- ^a Sled Configuration Details
- Minimum: Each MX740c is configured with 2x 85 W CPUs, 12x 16 GB DIMMs, 2x 2.5" SATA HDDs, 1x 25 Gb Mezz card, 1x H740 PERC
- Typical: Each MX740c is configured with 2x 145 W CPUs, 12x 32 GB DIMMs, 6x 1600 GB SAS SSD, 2x 25 GB Mezz cards, 1x H740P PERC, 1x Fiber-channel mini Mezz
- Each MX840c is configured with 4x 165 W CPUs, 48x 16 GB DIMMs, 8x 1600 GB NVMe SSD, 4x 25 GB Mezz cards, 2x Fiber-channel mini Mezz, 1x H740P PERC
- Feature-Rich: Each MX740c is configured with 2x 165 W CPUs, 24x 32 GB DIMMs, 6x 1600 GB NVMe SSD, 2x 25 GB Mezz cards, 1x H740P PERC
- Each MX5016s is configured with 16x 1600 GB SSD
- ^b LwA-UL is the upper limit sound power levels (LwA) calculated per section 4.4.1 of ISO 9296 (1988) and measured in accordance to ISO 7779 (2010)
- ^c Idle values are recorded when server is powered on with only the operating system running

9. ^dProminent tone: Criteria of D.6 and D.11 of ECMA-74 13th ed. (2015) are followed to determine if discrete tones are prominent. The acoustic transducers are at front and rear bystander positions. The system is placed in a 42U rack enclosure, 75 cm above the reflective floor.

Acoustical Baffle

The acoustical baffle is an optional add-on developed for MX7000 to provide reduced acoustics and improved sound quality for discerning customers. The baffle attaches to the rear of the chassis and provides attenuation of acoustics that are related to the rear fans (sled cooling). The baffle is designed for tool-less attachment. The baffle has minimal impact to chassis airflow, but is not recommended for higher wattage CPUs (>140 W TDP) in environments where ambient exceeds 35°C. Under these conditions, system performance may be impacted by the baffle.

 **CAUTION: The acoustical baffle is not designed to support the weight of cables, therefore no cables should be attached to, or rest on, the baffle.**

Sound Cap

Sound cap is a new custom thermal profile that is designed for customers that temporarily require reduced acoustics (reboot required) and are willing to accept some reduced performance. Sound cap works by applying a percentage power cap to the CPUs in the sled. Sound cap is only available on MX740c and MX840c sleds. Sound cap is not available as a chassis-level setting.

Acoustical Dependencies

The following statements reflect the most important system configuration dependencies that drive acoustical experience:

- Generally, the sled fans (rear fan modules) are the loudest component in the system, therefore reducing the thermal load on individual sleds is the most successful approach to reducing acoustics.
- For compute sled configurations (MX740c & MX840c), CPU thermal design power (TDP) and CPU case temperature requirements drive the cooling requirements of the sled for most workloads. Lower TDP processors that are coupled with workload scheduling may reduce acoustics.
- For IOM-A/B options, the best acoustics are achieved with 10G, 25 Gb modules, and fabric expander (MX7116n) modules, moderate acoustical experience can be achieved with the switching module (MX5108n), while the fabric switching module (MX9116n) has the loudest acoustical experience.
- For IOM-C options, SAS storage IOM (MX5000s) requires less cooling than the fiber channel module (MXG610s) thus lower acoustics.
- When sled slots are empty, blanks must be installed to achieve efficient cooling for sleds. Inefficient cooling induces loud acoustics

OpenManage systems management

The Dell EMC OpenManage systems management portfolio includes powerful hardware and software management tools and consoles. OpenManage simplifies the lifecycle of deploying, updating, monitoring and maintaining your Dell EMC PowerEdge servers.

Topics:

- [iDRAC with Lifecycle controller](#)
- [Agent-free management](#)
- [Agent-based management](#)
- [OpenManage console](#)
- [OpenManage systems management tools, utilities and protocols](#)
- [OpenManage Integrations](#)
- [OpenManage Connections](#)

iDRAC with Lifecycle controller

The integrated Dell Remote Access Controller (iDRAC) with Lifecycle Controller provides embedded management in every Dell EMC PowerEdge server. It provides functionality that helps you deploy, update, monitor, and maintain Dell EMC PowerEdge servers with or without a systems management software agent, and because it is embedded, iDRAC does not need an operating system or hypervisor to start working. This agent-free operation means that with OpenManage technology, your new server is ready to deploy a new operating system or accept new settings without installing extra software. And with powerful, easy-to-use, remote management and configuration options, iDRAC with Lifecycle Controller can alert you when an issue occurs, no matter where you are.

OpenManage Enterprise

OpenManage Enterprise simplifies, centralizes, and automates the full span of server life cycle management activities. It helps in discovery, configuration, deployment, updates, and remediation. These tasks are performed within a single console that unifies management of tower, rack, and modular platforms. OpenManage Enterprise helps in standardizing and supporting IT management policies and practices.

The OpenManage Enterprise console provides capabilities in the following areas:

- Reduced time and effort that is required to manage small and large-scale IT environments using HTML5 GUI.
- Simplified GUI workflow to provide a single management layer that unifies the management of PowerEdge tower, rack, and modular platforms.
- Packaged and delivered as virtual appliance and supports ESXi, Hyper-V, and KVM.
- Redesigned architecture on CentOS with the PostgreSQL database. Operating system and database license are no longer required.
- Centralized user management with role-based access control.
- Enabled with customer automation and solution integration using Northbound API.
- Enhanced policy-driven management.

OpenManage Enterprise features and benefits

Table 14. OpenManage Enterprise features and benefits

Feature	Description	Benefits
Packaged as a virtual appliance	OpenManage Enterprise is readily deployable as a virtual appliance in ESXi, Hyper-V, and KVM hypervisors.	Enables quick and easy deployment while gaining the benefits of virtualization,

Table 14. OpenManage Enterprise features and benefits (continued)

Feature	Description	Benefits
		including the migration of the application from one physical server to another.
Discovery and Inventory on the IPv4 or IPv6 network	Discovery and inventory data center hardware up to 5500 devices.	Provides data-center visibility for IT administrators. Improves security and provides a management solution for customers no longer using IPv4.
Automatic device onboarding	The "make ready for management" step is automated for iDRAC during the device discovery and inventory process.	Minimizes human errors during the process of discovery and setup. Helps in getting the environment ready for on-going management.
Device level role-based access control (RBAC)	The scope of IT administration responsibilities can be assigned to specific groups of devices.	Helps ensure IT Admin workload balance and skill-set alignment.
Server single sign-on	Use of credential caching is employed to avoid further password prompts when IT administration perform management tasks or launch the iDRAC virtual console.	Delivers fast and seamless management across the PowerEdge server lifecycle while enabling management from a single level.
Monitoring and Event Processing	Monitoring system health for discovered devices and create event policy to dispatch hardware events.	Provides IT admin real-time feedback for the health of their environment and enables automated responses through automation policy.
Comprehensive API	All product features are accessible from a new RESTful API that is conformant to the new popular Redfish standard.	Enables IT Admins to automate features and functions of OpenManage Enterprise through tools such as PowerShell and Python. It also allows them the flexibility of integrating OpenManage Enterprise into their existing IT processes.
Custom reports	IT administrators can use a powerful reporting engine to create customized reports.	Provides IT Admins rapid and easy access to specific device information.
Modern UI with elastic search capabilities	OpenManage Enterprise uses the modern HTML5 standard for its UI while enabling an elastic search engine that allows IT administrators to find anything in the console in a single search	Minimizes training time and maximizes efficiency by delivering quick results on searches involving devices, hardware and software inventory, features, functions, and so on. Also provides access from various browsers, and mobile devices.
Firmware compliance policies and updates	Firmware compliance policies allow IT administrators to establish one or more firmware baselines for groups of the PowerEdge servers in their environment. Selectively update firmware that is not compliant with the customer defined baselines.	Provides a simple method of enforcing minimum firmware levels for the PowerEdge servers that eliminates human error and minimizes downtime. Reinforces security by enabling customers to stay current with the firmware recommended by Dell EMC.
Configuration compliance policies	Configuration compliance policies allow IT Administrators to establish one or more configuration baselines for groups of the PowerEdge servers in their environment.	Enables a simple method to monitor the PowerEdge servers against an established baseline. Helps in enforcing security standards and ensuring management consistency across server BIOS, RAID, network, and the iDRAC settings.
Bare-metal server deployment	Automate server HW configuration deployment to a bare-metal PowerEdge	Saves time and relies less on local IT Administrators for provisioning and

Table 14. OpenManage Enterprise features and benefits (continued)

Feature	Description	Benefits
	server and kick off unattended operating system installation process.	configuring PowerEdge server hardware in a reliable and consistent manner.
Custom Template Edit	Edit the deployment template attributes by using: <ul style="list-style-type: none"> • Guided view • Advanced view 	Provides ability to control template editing using "guided view" for specific attributes, such as BIOS, boot sequence, and so on, to avoid human error. Granular control for sophisticated users using the "advanced view" template editing
Support for MX-series	Support for discovery, inventory, configuration compliance and remediation of MX7000 chassis-as a stand-alone chassis and as a lead chassis in a Multi Chassis Management (MCM) group.	Provides unified support for rack, blades, chassis from a single choice of console for our customer
Stateless Deployments and virtual Identity management	You can assign virtual identities to servers and perform stateless deployment.	Stateless deployments enable to create a server. environment that is dynamic and flexible by automatically generating and assigning virtual identity attributes to the I/O interfaces of a server
New Platform Support	Support for the latest 14th generation PowerEdge servers including the new blade servers of MX7000 chassis.	Provides unified support for multigenerational rack, blades, chassis from a single choice of console for our customer
OpenManage Enterprise RESTful API	REST API enhancements for network configuration, identity pool creation, device template configuration, and profile management	Helps customer script from higher-level frameworks to enable enhanced automation in a data center environment.
Configuration Remediation	Remediation of a server config that has drifted against a defined baseline.	Helps keep data center in compliance to configuration standards.

OpenManage Enterprise - Modular Edition

The Dell EMC OpenManage Enterprise-Modular Edition (OME-Modular Edition) helps to transform how you manage the infrastructure and enable the business quickly. It is a unified web/RESTful API interface that manages all the nodes including compute, storage, and networking. This helps reduce costs and consolidates multiple tools for ease of access and monitoring. The simplified administration helps deploy and monitor at scale, from one to many chassis, with the support for remote management. The Agile and intelligent automation drives faster hardware rollouts and reduces repetitive tasks for accelerated life cycle management.

It runs on the PowerEdge M9002m management module (MM) firmware. OME-Modular facilitates configuration and management of a stand-alone PowerEdge MX7000 chassis or group of MX7000 chassis(up to 20 or 160 servers) from one Graphical User Interface (GUI). You can use OME-Modular to deploy servers and update firmware. You can also manage the overall health of the chassis and the chassis components such as compute sleds, networks, input or output modules (IOMs), and storage devices. OME-Modular also facilitates hardware enabling:

- Connectivity of management network
- Discovery, inventory
- Monitoring and power control operations and thermal functions

You can use OME-Modular to manage key workloads on the MX7000 platforms.

- Large and unstructured data and analytics
- Hyper converged and traditional workloads
- Database workloads
- Software defined storage
- HPC and performance workloads

The lead chassis in the Multi Chassis Management (MCM) enables you to perform the following tasks:

- Manage servers across multiple MX chassis.
- Deploy or update servers from lead chassis without launching the member chassis web interface.
- Manage fabric switch engines in fabric mode using the OME-Modular web interface.
- Manage alert log and actions.
- Manage virtual MAC/WWN identity pools.
- Deploy compute sleds easily using server profiles and templates.

OME-Modular offers simple and static roles such as the chassis administrator, compute manager, fabric manager, storage manager, and viewer roles while, OpenManage Enterprise offers static and dynamic groups with role-based access control (RBAC).

Agent-free management

As Dell EMC PowerEdge servers have embedded server lifecycle management, there is no need to install an OpenManage systems management software agent into the operating system of a PowerEdge server. This greatly simplifies and streamlines the management footprint.

Agent-based management

Most systems management solutions require pieces of software, called agents, to be installed on each node in order to be managed within the IT environment. Additionally, the same agent is often used as a local interface into the hardware health and may be accessed remotely as a management interface, typically referred to as a one-to-one interface. For customers that continue to use agent-based solutions, Dell EMC provides OpenManage Server Administrator.

OpenManage console

The central console in a systems management solution is often referred to as the one-to-many console. The central console provides a rapid view and insight into the overall health of all systems in the IT environment. The Dell EMC systems management portfolio includes several powerful consoles, depending upon your needs, including the following:

- **OpenManage Enterprise-Modular Edition** - OpenManage Enterprise - Modular Edition (OME - Modular) is the systems management software that runs on the Dell EMC PowerEdge MX-series chassis. OME-Modular shares a common code base with OpenManage Enterprise, delivering an agile, reliable and comprehensive platform to manage PowerEdge MX servers, storage and networking. OME-Modular boosts services delivery and restores IT agility for growing businesses. A unified web and RESTful API interface manages all nodes including compute, storage, and networking. This helps reduce costs, learning curve, and consolidates multiple tools for ease of access and monitoring. Simplified administration helps deploy and monitor at scale, from one to many chassis, with support for remote management. The agile and intelligent automation of OME - Modular helps enable faster hardware deployments and reduces repetitive tasks for accelerated life cycle management. Some of the key features of OME-Modular include:
 - Shares the same code base as OpenManage Enterprise making it easy for customers to move from one console to the other – or utilize both
 - Capability to manage servers across multiple MX-series chassis
 - Fast deployment with automatic chassis and node discovery
 - Comprehensive RESTful API to automate multiple tasks and integrate with third-party tools
 - Touchscreen display and OpenManage Mobile support for easy access
 - Expanded management across all PowerEdge servers with OpenManage Enterprise
 - Easily extend management into VMware vCenter and Microsoft System Center environments as well as others including Ansible, with Dell EMC Connections and Integrations
- **OpenManage Power Center**—Dell EMC's power management solution, the Dell EMC OpenManage Power Center (OMPC) management console, provides increased visibility to power consumption, anomalies, and utilization through fine-grained instrumentation. This enables increased control, improved rack density, faster response times, greater accuracy, and broader decision-making intelligence than would otherwise be possible. When used with a suitably licensed PowerEdge server (with a Dell iDRAC Enterprise license), OMPC leverages Intel Node Manager technology for platform-level power reporting and capping of Intel chipsets. Power Center then communicates with iDRAC to provide node, rack, row or data-center level aggregation of power-management data, as well as execution of control policy — making it easy for IT professionals to identify areas to gain efficiencies and cut wasteful costs.

OpenManage systems management tools, utilities and protocols

OpenManage systems management tools and utilities consist of the following:

- **Dell EMC Repository Manager:** The Dell EMC Repository Manager (RM) is a stand-alone GUI-based productivity tool that helps simplify the process of managing downloads and baseline BIOS, firmware and driver updates. Repository Manager can create deployment disks as well as create and manage customized repositories.
- **Dell EMC Update Packages:** The Dell EMC Update Packages (DUP) is a self-contained executable in a standard package format that updates a software element on a Dell EMC server such as the BIOS, a driver, firmware and other software updates.
- **Dell EMC OpenManage Deployment Toolkit:** The Dell EMC OpenManage Deployment Toolkit (DTK) is a CLI-based tool that includes a set of utilities for configuring and deploying Dell EMC PowerEdge systems, and can be used to build scripted, unattended OS installations to deploy large numbers of servers in a reliable fashion.
- **RACADM:** The RACADM command-line utility provides a scriptable interface that allows you to locally or remotely configure iDRAC7.
- **IPMITool:** IPMITool includes scriptable console application programs used to control and manage remote systems using the IPMI version 1.5 and later protocol.
- **Web Services for Management (WSMAN):** WSMAN is a SOAP-XML-based protocol for exchanging system management information. Dell EMC's implementation provides remote management capabilities through a secure and standards-based Web Services-Management (WS-MAN) interface to PowerEdge servers and blade server node chassis.

OpenManage Integrations

OpenManage provides integration with several leading third-party consoles, including:

- Dell EMC Deployment Pack for Microsoft System Center Configuration Manager
- Dell EMC PRO Management Pack for Microsoft System Center Virtual Machine Manager (SCVMM)
- Dell EMC Management Pack Suite for Microsoft System Center Operations Manager
- Lifecycle Controller Integration for Microsoft System Center Virtual Machine Manager
- Lifecycle Controller Integration for Microsoft System Center Configuration Manager
- OpenManage Integration for VMware vCenter
- BMC Software
- BMC TrueSight Server Automation
- BMC TrueSight Operations Management
- OpenManage Ansible Modules

OpenManage Connections

With OpenManage Connections, you can extract rich information about the health and behavior of Dell EMC servers and storage to more effectively monitor, manage and troubleshoot Dell EMC systems in multi-vendor hardware, OS, hypervisor and Open Source environments.

- Dell EMC OpenManage Micro Focus Operations Manager|Operations Connector
- IBM
 - OpenManage Connection for IBM Tivoli Netcool/OMNIBUS (ITNO)
 - OpenManage Connection for IBM Tivoli Network Manager (ITNM) IP Edition
- Nagios
 - OpenManage Plug-in for Nagios Core
 - OpenManage Plug-in for Nagios XI

Appendix A. Additional specifications

The following sections contain information about additional system specifications.

Topics:

- Population rules
- PSU redundancy and population rules
- Chassis dimensions
- Chassis weight
- Fan specifications
- PSU specifications
- Ports and connectors specifications
- PowerEdge MX modules ports and connectors
- Video specifications
- Environmental specifications

Population rules

System modules must be populated as described in the following table:

Table 15. MX7000 population rules




Category	Maximum population
Blanks	All the empty slots in the MX7000 enclosure must be populated with blanks (Sled, IOM, EC, and PSU). This is required for proper cooling of the enclosure and components.
Fans	All system fans must be populated in the enclosure.
Power supply units	The number of power supply units that are required is dependent upon system configuration and redundancy mode, and the minimum recommended quantity is two. The six PSUs are organized into two groups: Grid A consists of PSUs 1, 2, 3, and Grid B consists of PSUs 4, 5, 6. It is recommended to populate the PSUs in the following order: 1, 4, 2, 5, 3, 6 where an equal number of PSUs on each grid is optimized for Grid Redundancy. PSU redundancy and No Redundancy options do not have any PSU population requirements.
Power cable	One C21/C20 power cable must be connected to the C22 plug corresponding to each populated PSU.
Management module	A Management module must be present to control and manage the enclosure.  NOTE: If a single Management module crashes, the system functions normally.  NOTE: The enclosure cannot be managed or controlled until the Management module is replaced.
Control panel	The right control panel and one of the left control panel configurations (LCD or LED) must be present on the MX7000 enclosure.
Compute sled	Up to eight single-width or four double-width sleds or a combination can be populated. The double-width sleds must be in slots 1, 3, 5, 7 due to the enclosures design.
Storage Sled	Up to seven Storage sleds can be populated in the enclosure.  NOTE: One compute node must be present and it must be mapped to a storage node. One Fabric-C SAS IOM must be present and powered ON.
I/O Module	Only Brocade and SAS IOM are supported in Fabric-C. Only one type of IOM can be offered in Fabric-C (Fibre Channel or SAS IOM, not mixed).

Table 15. MX7000 population rules (continued)

Category	Maximum population
	Only one type of switch can be offered in Fabric-B (HPCC or Ethernet).
	Two Fabric-C SAS IOMs must be installed if the enclosure contains a Storage Node.
	Mix Speed of pass-through in same fabric is not enabled.
Mezzanine cards	If the enclosure contains a storage node, Fabric-C MiniMezzanine card (HBA330 or Jumbo PERC) must be installed in one compute node.
	Dual Port or quad port mezzanine cards must be installed for redundant IOM/Pass-through configurations.
	The second processor must be installed on the compute node to support Fabric-B Mezzanine / IOM and Fabric-C Mezzanine / IOM.

PSU redundancy and population rules

The number of PSUs required depends on the enclosure configuration and redundancy required. The minimum requirement is two PSUs. The enclosure supports one of the following redundancy modes:

- No redundancy:** This mode distributes the enclosure power load across all PSUs. There are no specific PSU population requirements for No redundancy. The intent of this mode is to have the highest possible limit for power enablement of devices that are added to the enclosure. If there are single or multiple PSU failures, then the enclosure limits performance to operate within the power capabilities of the remaining PSUs.
- Grid redundancy:** This mode distributes the enclosure power load across all PSUs. The six PSUs are organized into two groups: Grid A consisting of PSUs 1, 2, 3, and Grid B consists of PSUs 4, 5, 6. For grid redundancy, PSUs should be populated in the following order: 1, 4, 2, 5, 3, 6. The grid with the largest PSU capacity determines the limit for power enablement of devices that are added to the enclosure. If there is a grid or PSU failure, then the enclosure power is distributed among the remaining PSUs with the intent that a single healthy grid will continue to provide power to all the components in the system which can experience varied performance depending on the workload.
- PSU redundancy:** This mode distributes the enclosure power load across all PSUs. There are no specific PSU population requirements for PSU redundancy. PSU redundancy is optimized for a population of six PSUs, and the enclosure limits the power enablement of devices to fit within five PSUs. If there is a single PSU failure, then the enclosure power is distributed among the remaining PSUs to provide power to all the components in the system which can experience varied performance depending on the workload. If there are fewer than six PSUs, then the enclosure limits the power enablement of devices to fit within all populated PSUs. If there is a single PSU failure, then the enclosure limits performance to operate within the power capabilities of the remaining PSUs.

Table 16. PSU population rules

PSU count	Population order
2	1, 4 (Optimized for Grid Redundancy 1+1, and Hot Spare)
3	1, 4, 2
4	1, 4, 2, 5 (Optimized for Grid Redundancy 2+2, and Hot Spare)
5	1, 4, 2, 5, 3
6	1, 4, 2, 5, 3, 6 (Optimized for Grid Redundancy 3+3, PSU Redundancy 5+1, and Hot Spare)

Hot Spare: The MX7000 PSUs support the Hot Spare feature with three PSU pairs. This feature enables a PSU pair to have one active PSU and one PSU in sleep mode while the enclosure power consumption is low, and the three PSU pairs meet all the power requirements for the enclosure. This enables efficient power utilization when the overall enclosure power requirement is low. The partner PSU wakes the paired PSU from sleep mode by sending a WAKE signal when the enclosure power requirement increases. The PSU pairs for MX7000 are—1 & 4, 2 & 5, and 3 & 6.

Chassis dimensions




Figure 8. Dimensions of the PowerEdge MX7000

Table 17. Dimensions of the PowerEdge MX7000

Description	Dimension
Xa	482 mm (18.98 inches)
Xb	445 mm (17.52 inches)
Y	307.4 mm (12.11 inches)
Zb	811.6 mm (31.96 inches)
Zc	816.6 mm (32.15 inches)

Chassis weight

Table 18. Chassis weight

Enclosure	Minimum weight	Maximum weight (fully populated)
PowerEdge MX7000	82 kg (180 lbs)	182 kg (400 lbs)

Fan specifications

The PowerEdge MX7000 enclosure supports four front accessible hot-swap cooling fans and five rear accessible hot-swap cooling fans. The cooling fan assembly ensures that the key components of the server such as the sleds, Fabrics, and I/O modules get adequate air circulation to keep them cool. A cooling fan failure can result in overheating and may lead to damage.

Table 19. Supported fans

Fan location	Front	Rear
Size	60 mm	80 mm
Number of fans	4	5
Redundancy	3+1	4+1


PSU specifications


The PowerEdge MX7000 enclosure supports up to six AC power supply units (PSUs).


Table 20. PSU specifications

Description	Specification
PSU	6 x 3000 W AC
Class	Platinum
Heat dissipation (maximum)	1205 BTU/hr
Frequency	50/60 Hz
Voltage	100–240 V AC, autoranging
Current	6 x 16 A
Inrush current	Maximum 40 A per power supply for 10 ms or less
	Maximum 50 A per power supply for 1.2 ms or less
Connector	C22 AC power connector

 **CAUTION:** Mixed high line and low line AC inputs in the same enclosure is not supported.

 **NOTE:** Heat dissipation is calculated using the PSU wattage rating.

 **NOTE:** In an ideal input voltage condition and over the entire enclosure's ambient operating range, the inrush current may reach 120 A per power supply for 10 ms or less.

 **NOTE:** This enclosure is also designed to connect to the IT power enclosures with a phase-to-phase voltage not exceeding 240 V.


Ports and connectors specifications

USB ports

The PowerEdge MX7000 enclosure supports two Type A, USB 2.0 ports on the front panel.

Mini DisplayPort


The PowerEdge MX7000 enclosure supports one Mini DisplayPort (mini DP) on the front panel.

 **NOTE:** You must use a mini DP dongle to connect the enclosure to a VGA display.

PowerEdge MX modules ports and connectors


PowerEdge MX740c

Table 21. PowerEdge MX740c externally accessible connectors

Connector	Description
USB ports	<ul style="list-style-type: none">• One USB 3.0-compliant port on the front of the sled.• One USB 3.0-compliant internal port.• One micro USB 2.0-compliant port for iDRAC Direct on the front of the sled. <p> NOTE: The micro USB 2.0-compliant port on the front of the sled can only be used as an iDRAC Direct port.</p>

PowerEdge MX840c

Table 22. PowerEdge MX840c externally accessible connectors

Connector	Description
USB ports	<ul style="list-style-type: none">• One USB 3.0-compliant port on the front of the sled.• One USB 3.0-compliant port internal port.• One micro USB 2.0-compliant port for iDRAC Direct on the front of the sled. <p> NOTE: The micro USB 2.0-compliant port on the front of the sled can only be used as an iDRAC Direct port.</p>

MX7116n Fabric Expander Module

Table 23. MX7116n Fabric Expander Module externally accessible connectors

Connector	Description
Externally accessible connectors	<ul style="list-style-type: none">• 2 QSFP28-DD connections to the MX7116n

MX9116n Fabric Switching Engine

Table 24. MX9116n Fabric Switching Engine externally accessible connectors

Connector	Description
Externally accessible connectors	<ul style="list-style-type: none">• 12 QSFP28-DD ports that can be configured as:<ul style="list-style-type: none">○ 2 x 40 GbE or 2 x 100 GbE ports for uplinks○ 8 x 10 GbE or 8 x 25 GbE ports for rack servers• 2 QSFP28 uplink ports that can be configured as:<ul style="list-style-type: none">○ 1 x 40 GbE○ 1 x 100 GbE○ 2 x 50 GbE○ 4 x 10 GbE○ 4 x 25 GbE• 2 QSFP28 unified ports that can be configured as:<ul style="list-style-type: none">○ 1 x 40 GbE

Table 24. MX9116n Fabric Switching Engine externally accessible connectors

Connector	Description
	<ul style="list-style-type: none"> ○ 1 x 100 GbE ○ 2 x 50 GbE ○ 4 x 10 GbE ○ 4 x 25 GbE ○ 8 x 8/ 16/ 32 GbE Fibre Channel mode

MX9002m Management Module

Table 25. MX9002m Management Module externally accessible connectors

Connector	Description
Externally accessible connectors	<ul style="list-style-type: none"> ● Two x 1G-BaseT Ethernet ports ● One x Micro-B USB port

MX5108n Ethernet Switch

Table 26. MX5108n Ethernet Switch externally accessible connectors

Connector	Description
Externally accessible connectors	<ul style="list-style-type: none"> ● 2 x 100 GbE QSFP28 uplink ports ● 1 x 40 GbE QSFP+ uplink port ● 4 x 10GBASE-T uplink ports ● USB Serial and USB Flash ports

MXG610s Fibre Channel Switch

Table 27. MXG610s Fibre Channel Switch externally accessible connectors

Connector	Description
USB port	One micro USB 2.0-compliant port on the front of the sled.
Fibre Channel transceiver	16 external ports supporting 8/ 16/ 32 Gbps speeds using 8 SFPs and 2 QSFPs.

PowerEdge MX 10GBASE-T Ethernet Pass-Through Module

Table 28. PowerEdge MX 10GBASE-T Ethernet Pass-Through Module externally accessible connectors

Connector	Description
Fibre Channel transceiver	16 external ports supporting 10 GbE connections

PowerEdge MX 25 Gb Ethernet Pass-Through Module

Table 29. PowerEdge MX 25 Gb Ethernet Pass-Through Module externally accessible connectors

Connector	Description
Fibre Channel transceiver	16 external ports supporting 25 GbE connections

Video specifications

The management module supports an integrated Matrox G200eW3 graphics controller with a 16 MB video frame buffer.

Table 30. Supported video resolution options

Resolution	Refresh rate Hz	Color depth (bits)
1024 x 768	60	8, 16, 32
1280 x 800	60	8, 16, 32
1280 x 1024	60	8, 16, 32
1360 x 768	60	8, 16, 32
1440 x 900	60	8, 16, 32
1600 x 900	60	8, 16, 32
1600 x 1200	60	8, 16, 32
1680 x 1050	60	8, 16, 32
1920 x 1080	60	8, 16, 32
1920 x 1200	60	8, 16, 32

NOTE: 1920 x 1080 and 1920 x 1200 resolutions are only supported in reduced blanking mode.

Environmental specifications

NOTE: For additional information about environmental measurements for specific system configurations, see www.dell.com/poweredgemanuals.

Table 31. Temperature specifications

Temperature	Specifications
Storage	-40°C to 65°C (-40°F to 149°F)
Maximum temperature gradient (Operating and storage)	20°C/h (36°F/h)

Table 32. Relative humidity specifications

Relative humidity	Specifications
Storage	5% to 95% RH with 33°C (91°F) maximum dew point. Atmosphere must be noncondensing always.
Operating	10% to 80% RH with 29°C (84.2°F) maximum dew point.

Table 33. Maximum vibration specifications

Maximum vibration	Specifications
Operating	0.26 G _{rms} at 5 Hz to 350 Hz (all axis)
Storage	1.88 G _{rms} at 10 Hz to 500 Hz (vertical axis)

Table 34. Maximum shock pulse specifications

Maximum shock pulse	Specifications
Operating	Shock pulses in the positive and negative x, y, and z axis of 6 G for up to 11 ms.

Table 34. Maximum shock pulse specifications (continued)

Maximum shock pulse	Specifications
Storage	Shock pulses in the positive z axis of 71 G for up to 2 ms. Shock pulses in the positive and negative x and y axis of 20 G for up to 7 ms.

Table 35. Maximum altitude specifications

Maximum altitude	Specifications
Operating	3048 m (10,000 ft)
Storage	12,000 m (39,370 ft)

Table 36. Operating temperature derating specification

Operating temperature derating	Specifications
Up to 35°C (95°F)	Maximum temperature reduces by 1°C/300 m (1°F/547 ft), above 950 m (3,117 ft).
35°C to 40°C (95°F to 104°F)	Maximum temperature reduces by 1°C/175 m (1°F/319 ft), above 950 m (3,117 ft).
40°C to 45°C (104°F to 113°F)	Maximum temperature reduces by 1°C/125 m (1°F/228 ft), above 950 m (3,117 ft).

Standard operating temperature

Table 37. Standard operating temperature specifications

Standard operating temperature	Specifications
Continuous operation (For altitude less than 950 m or 3117 ft)	10°C to 35°C (50°F to 95°F) with no direct sunlight on the equipment.

Expanded operating temperature

Table 38. Expanded operating temperature specifications

Expanded operating temperature	Specifications
Continuous operation	5°C to 40°C at 5% to 85% RH with 29°C dew point. i NOTE: Outside the standard operating temperature (10°C to 35°C), the system can operate continuously in temperatures as low as 5°C and as high as 40°C. For temperatures 35°C to 40°C, derate maximum allowable temperature by 1°C per 175 m (1°F per 319 ft) above 950 m (3,117 ft).
≤ 1% of annual operating hours	–5°C to 45°C at 5% to 90% RH with 29°C dew point. i NOTE: Outside the standard operating temperature (10°C–35°C), the system can operate down to –5°C or up to 45°C for a maximum of 1% of its annual operating hours. For temperatures between 40°C and 45°C, derate maximum allowable temperature by 1°C per 125 m (1°F per 228 ft) above 950 m (3,117 ft).

i **NOTE:** When operating in the expanded temperature range, the performance of the system may be impacted.

i **NOTE:** When operating in the expanded temperature range, ambient temperature warnings may be reported on the LCD panel and in the System Event Log.

Expanded operating range

- The operating temperature is specified for a maximum altitude of 950 m for expanded operating range.
- Do not perform cold start at 5°C or lower, due to hard drive constraints.
- Redundant power supplies are required.

Expanded operating temperature restrictions

For more information about the expanded operating temperature restrictions, see the Installation and Service Manual for the PowerEdge MX sleds at www.dell.com/poweredgemanuals.

Table 39. Expanded operating temperature restrictions

System	C30	C35	C40E45
Dell EMC PowerEdge MX7000 including the fans, Management Module, and PSUs	No restrictions	No restrictions	No restrictions
Fabrics A and B modules	No restrictions	No restrictions	The MX9116n is not supported.
Fabric C I/O modules	No restrictions	No restrictions	No restrictions

Particulate and gaseous contamination specifications


The following table defines the limitations that help avoid any damages to the IT equipment and/or failure from particulates and gaseous contamination: If the levels of particulates or gaseous pollution exceed the specified limitations and result in equipment damage or failure, you may need to rectify the environmental conditions. Remediation of environmental conditions is the responsibility of the customer.

Table 40. Particulate contamination specifications

Particulate contamination	Specifications
Air Filtration	Data center air filtration defines that, ISO Class 8 per ISO 14644-1 with a 95% upper confidence limit. <i>i</i> NOTE: This condition applies to data center environments only. Air filtration requirements do not apply to IT equipment designed to be used outside a data center, in environments such as an office or factory floor. <i>i</i> NOTE: Air entering the data center must have MERV11 or MERV13 filtration.
Conductive dust	Air must be free of conductive dust, zinc whiskers, or other conductive particles. <i>i</i> NOTE: This condition applies to data center and nondata center environments.
Corrosive dust	<ul style="list-style-type: none"> • Air must be free of corrosive dust. • Residual dust present in the air must have a deliquescent point less than 60% relative humidity. <i>i</i> NOTE: This condition applies to data center and nondata center environments.

Table 41. Gaseous contamination specifications

Gaseous contamination	Specifications
Copper Coupon Corrosion	<300 Å/month per Class G1 defines that, ANSI/ISA71.04-1985.
Silver Coupon Corrosion	<200 Å/month defines that, AHSRAE TC9.9.

 **NOTE:** Maximum corrosive contaminant levels measured at $\leq 50\%$ relative humidity.

Appendix B. Standards compliance

Table 42. Industry standard documents

Standard	URL for information and specifications
ACPI Advance Configuration and Power Interface Specification, v2.0c	acpi.info
Ethernet IEEE 802.3-2005	standards.ieee.org/getieee802/802.3.html
HDG Hardware Design Guide Version 3.0 for Microsoft Windows Server	microsoft.com/whdc/system/platform/pcdesign/designguide/serverdg.msp
IPMI Intelligent Platform Management Interface, v2.0	intel.com/design/servers/ipmi
DDR4 Memory DDR4 SDRAM Specification	jedec.org/standards-documents/docs/jesd79-4.pdf
PCI Express PCI Express Base Specification Rev. 2.0 and 3.0	pcisig.com/specifications/pciexpress
PMBus Power System Management Protocol Specification, v1.2	pmbus.info/specs.html
SAS Serial Attached SCSI, v1.1	t10.org
SATA Serial ATA Rev. 2.6; SATA II, SATA 1.0a Extensions, Rev. 1.2	sata-io.org
SMBIOS System Management BIOS Reference Specification, v2.7	dmtf.org/standards/smbios
TPM Trusted Platform Module Specification, v1.2 and v2.0	trustedcomputinggroup.org
UEFI Unified Extensible Firmware Interface Specification, v2.1	uefi.org/specifications
USB Universal Serial Bus Specification, Rev. 2.0	usb.org/developers/docs

Appendix C. Additional resources

Table 43. Additional resources

Resource	Description of contents	Location
Installation and Service Manual	<p>This manual, available in PDF format, provides the following information:</p> <ul style="list-style-type: none"> • Chassis features • System Setup program • System messages • System codes and indicators • System BIOS • Remove and replace procedures • Troubleshooting • Diagnostics • Jumpers and connectors 	Dell.com/Support/Manuals
Getting Started Guide	<p>This guide ships with the system, and is also available in PDF format. This guide provides the following information:</p> <ul style="list-style-type: none"> • Initial setup steps • Key system features • Technical specifications 	Dell.com/Support/Manuals
Rack Installation Instructions	<p>This document ships with the rack kits, and provides instructions for installing a server in a rack.</p>	Dell.com/Support/Manuals
Information Update	<p>This document ships with the system, is also available in PDF format online, and provides information on system updates.</p>	Dell.com/Support/Manuals
Energy Smart Solution Advisor (ESSA)	<p>The Dell online ESSA enables easier and more meaningful estimates to help you determine the most efficient configuration possible. Use ESSA to calculate the power consumption of your hardware, power infrastructure, and storage.</p>	Dell.com/calc

Appendix D. Support and deployment services

Topics:

- Dell EMC ProDeploy Enterprise Suite
- Deployment services
- Dell EMC Remote Consulting Services
- Dell EMC Data Migration Service
- ProSupport Enterprise Suite
- ProSupport Plus
- ProSupport
- ProSupport One for Data Center
- Support Technologies
- Additional professional services
- Dell Education Services
- Dell EMC Global Infrastructure Consulting Services
- Dell EMC Managed Services

Dell EMC ProDeploy Enterprise Suite

ProDeploy Enterprise Suite gets your server out of the box and into optimized production—fast. Our elite deployment engineers with broad and deep experience utilizing best-in-class processes along with our established global scale can help you around the clock and around the globe. From simple to the most complex server installations and software integration, we take the guess work and risk out of deploying your new server technology.

		Basic Deployment	ProDeploy	ProDeploy Plus
Pre-deployment	Single point of contact for project management		•	In-region
	Site readiness review		•	•
	Implementation planning		•	•
	Technology Service Manager (TSM) engagement for ProSupport Plus entitled devices			•
Deployment	Deployment service hours	Business hours	24x7	24x7
	Onsite hardware installation*	•	•	•
	Packaging materials disposal	•	•	•
	Install and configure system software		•	Onsite
Post-deployment	Project documentation with knowledge transfer		•	•
	Deployment verification		•	•
	Configuration data transfer to Dell EMC technical support		•	•
	30-days of post-deployment configuration assistance			•
	Training credits for Dell EMC Education Services			•

Figure 9. ProDeploy Enterprise Suite capabilities

NOTE: Hardware installation not applicable on selected software products.

Dell EMC ProDeploy Plus

From beginning to end, ProDeploy Plus provides the skill and scale needed to successfully execute demanding deployments in today's complex IT environments. Certified Dell EMC experts start with extensive environmental assessments and detailed migration planning and recommendations. Software installation includes set up of most versions of Dell EMC SupportAssist and

OpenManage system management utilities. Post-deployment configuration assistance, testing, and product orientation services are also available.

Dell EMC ProDeploy

ProDeploy provides full service installation and configuration of both server hardware and system software by certified deployment engineers including set up of leading operating systems and hypervisors as well as most versions of Dell EMC SupportAssist and OpenManage system management utilities. To prepare for the deployment, we conduct a site readiness review and implementation planning exercise. System testing, validation, and full project documentation with knowledge transfer complete the process.

Dell EMC Basic Deployment

Basic Deployment delivers worry-free professional installation by experienced technicians who know Dell EMC servers inside and out.

Dell EMC Residency Services

Residency Services helps customers transition to new capabilities quickly with the assistance of on-site or remote Dell EMC experts whose priorities and time you control. Residency experts can provide post implementation management and knowledge transfer related to a new technology acquisition or day-to-day operational management of the IT infrastructure.

Deployment services

Deployment services details and exceptions can be found in service description documents at the Enterprise Configuration and Deployment page on Dell.com.

Dell EMC Remote Consulting Services

When you are in the final stages of your PowerEdge server implementation, you can rely on Dell EMC Remote Consulting Services and our certified technical experts to help you optimize your configuration with best practices for your software, virtualization, server, storage, networking, and systems management.

Dell EMC Data Migration Service

Protect your business and data with our single point of contact to manage your data migration project. Your project manager will work with our experienced team of experts to create a plan using industry-leading tools and proven processes based on global best practices to migrate your existing files and data so your business system get up and running quickly and smoothly.

ProSupport Enterprise Suite

With Dell EMC ProSupport Services, we can help you keep your operation running smoothly, so you can focus on running your business. We will help you maintain peak performance and availability of your most essential workloads. Dell EMC ProSupport is a suite of support services that enable you to build the solution that is right for your organization. Choose support models based on how you use technology and where you want to allocate resources. From the desktop to the data center, address everyday IT challenges, such as unplanned downtime, mission-critical needs, data and asset protection, support planning, resource allocation, software application management and more. Optimize your IT resources by choosing the right support model.

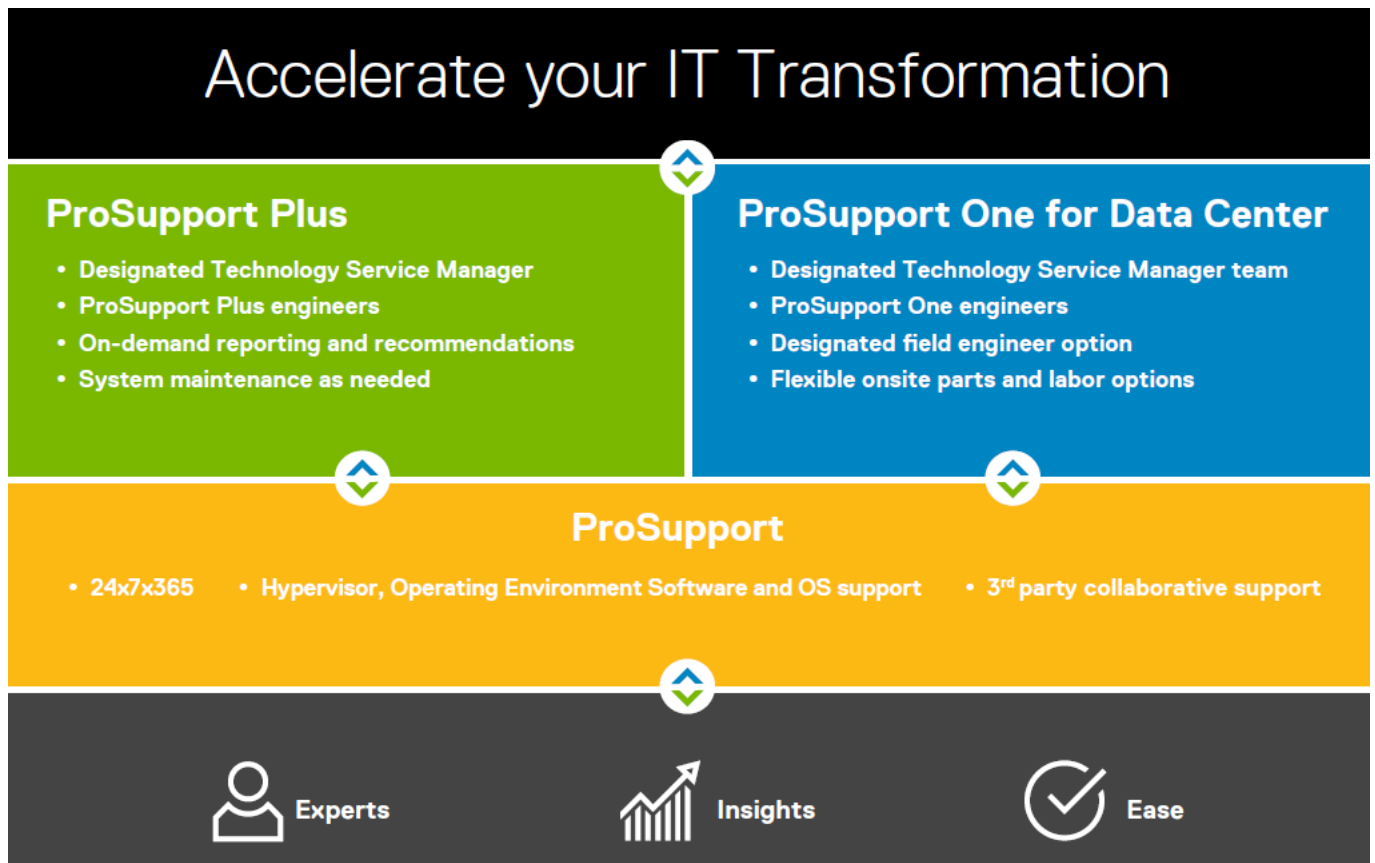


Figure 10. ProSupport Enterprise Suite

ProSupport Plus

When you purchase PowerEdge servers, we recommend ProSupport Plus, our proactive and preventative support, for business-critical systems. ProSupport Plus provides all the benefits of ProSupport, plus the following:

- An assigned Services Account Manager (SAM) who knows your business and your environment
- Access to senior ProSupport engineers for faster issue resolution
- Personalized, preventive recommendations based on analysis of support trends and best practices from across the Dell EMC customer base to reduce support issues and improve performance
- Predictive analysis for issue prevention and optimization enabled by SupportAssist
- Proactive monitoring, issue detection, notification and automated case creation for accelerated issue resolution enabled by SupportAssist
- On-demand reporting and analytics-based recommendations enabled by SupportAssist and TechDirect

ProSupport

Our ProSupport service offers highly trained experts around the clock and around the globe to address your IT needs. We will help you minimize disruptions and maximize availability of your PowerEdge server workloads with:

- 24x7x365 access to certified hardware and software experts
- Collaborative 3rd party support
- Hypervisor and OS support
- Consistent level of support available for Dell EMC hardware, software and solutions
- Onsite parts and labor response options including next business day or four-hour mission critical

ProSupport One for Data Center

ProSupport One for Data Center offers flexible site-wide support for large and distributed data centers with more than 1,000 assets. This offering is built on standard ProSupport components that leverage our global scale but are tailored to your company's needs. While not for everyone, it offers a truly unique solution for Dell EMC's largest customers with the most complex environments.

- Team of assigned Services Account Managers (SAM) with remote, on-site options
- Assigned ProSupport One technical and field engineers who are trained on your environment and configurations
- On-demand reporting and analytics-based recommendations enabled by SupportAssist and TechDirect
- Flexible on-site support and parts options that fit your operational model
- A tailored support plan and training for your operations staff

	ProSupport	ProSupport Plus	ProSupport One for Data Center
Remote technical support	24x7	24x7	24x7
Parts and labor response options	Next business day or Mission Critical	Next business day or Mission Critical	Flexible
Automated issue detection and case creation	•	•	•
Self-service case initiation and management	•	•	•
Hypervisor and OS support	•	•	•
Priority access to specialized support experts		•	•
Designated Technology Service Manager		•	•
Personalized assessments and recommendations		•	•
On-demand support and utilization reports		•	•
Systems Maintenance guidance		Semiannual	Optional
Designated technical and field support teams			•

Figure 11. Enterprise Support feature comparison

Support Technologies

Powering your support experience with predictive, data-driven technologies.

SupportAssist

The best time to solve a problem is before it happens. The automated proactive and predictive technology SupportAssist* helps reduce steps and time to resolution, often detecting issues before they become a crisis. Benefits include:

- Value - SupportAssist is available to all customers at no additional charge.
- Improve productivity - replace manual, high-effort routines with automated support.
- Accelerate time to resolution - receive issue alerts, automatic case creation and proactive contact from Dell EMC experts.
- Gain insight and control - optimize enterprise devices with on-demand ProSupport Plus reporting in TechDirect and get predictive issue detection before the problem starts.

SupportAssist is included with all support plans but features vary based on service level agreement.

	Basic Hardware Warranty	ProSupport	ProSupport Plus
Automated issue detection and system state information collection	•	•	•
Proactive, automated case creation and notification		•	•
Predictive issue detection for failure prevention			•
Recommendation reporting available on-demand in TechDirect			•

Figure 12. SupportAssist model

Get started at Dell.com/SupportAssist

TechDirect

Boost your IT teams productivity when supporting Dell EMC systems. With over 1.4 million self-dispatches processed each year, TechDirect has proven its effectiveness as a support tool. You can:

- Self-dispatch replacement parts
- Request technical support
- Integrate APIs into your help desk

Or, access all your Dell EMC certification and authorization needs. Train your staff on Dell EMC products as TechDirect allows you to:

- Download study guides
- Schedule certification and authorization exams
- View transcripts of completed courses and exams

Register at techdirect.dell.com

Additional professional services

Dell Education Services

Dell Education Services offers the PowerEdge server training courses designed to help you achieve more with your hardware investment. The curriculum is designed in conjunction with the server development team, as well as Dell EMC's technical support team, to ensure that the training delivers the information and practical, hands-on skills you and your team need to confidently manage and maintain your Dell EMC server solution. To learn more or register for a class today, visit LearnDell.com/Server.

Dell EMC Global Infrastructure Consulting Services

Dell EMC Global Infrastructure Consulting Services use skilled solution architects, innovative tools, automated analysis and Dell EMC's intellectual property to give rapid insight into the root causes of unnecessary complexity. We seek better answers than traditional service models, and our strategy is to help quickly identify high-impact, short-duration projects that deliver return on investment (ROI) and free up resources. The results are practical, action-oriented plans with specific, predictable, measurable outcomes. From data center optimization to server virtualization to systems management, our consulting services can help build a more efficient enterprise.

Dell EMC Managed Services

Dell EMC Managed Services are a modular set of lifecycle services designed to help you automate and centrally configure, deploy, and manage your day-to-day data center operations. These services extend your existing on-premise IT infrastructure with off-premise cloud services designed to better address challenges with mobility, highly distributed organizations, security, compliance, business continuity, and disaster preparedness.