Dell PowerEdge IR7000, M7701, and M7725

Technical Guide

Regulatory Model: E119S Series
Regulatory Type: E119S001
November 2025
Rev. A02

DELL Technologies

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.

Contents

1

PowerEdge IR7044 and IR7050 configurations and features

The IR7044 and IR7050 are compatible with multiple generations of compute tray enclosures (servers), ensuring that IT investments designed for use with Open Rack Ver.3 (ORv3) base rack can be reused multiple times.

Rack-Scale Platforms and Infrastructure

- Rack: Open Rack Ver.3 Base specification 44 OU and 50 OU which supports PowerEdge M7725

  NOTE: Note: 1 OU = 48 mm (1.88 inch) height, and 538.98 mm (21.22 inch) width.

- Direct liquid cooling (DLC) Manifolds: Dell Designed 21-inch UQDB06 based blind-mate DLC manifolds
- Busbar: ORv3 1400A
- Power Shelf: ORv3 33 kW
- In-Row and In-Rack CDUs
- Management Switch, for example: Dell S5248 10/25GbE
- Fabric Switch: Dell Ethernet and NVIDIA InfiniBand

Figure 1. IR7044 and IR7050

Topics:

- Blind mating for power busbar and Direct Liquid Cooling Universal Quick Disconnect (DLC UQD)

Front, rear, and side view of the IR7044 and IR7050

NOTE: The IR7050 front, rear, and side views are similar to IR7044. The only change in the height of IR7050 is 50 OU.

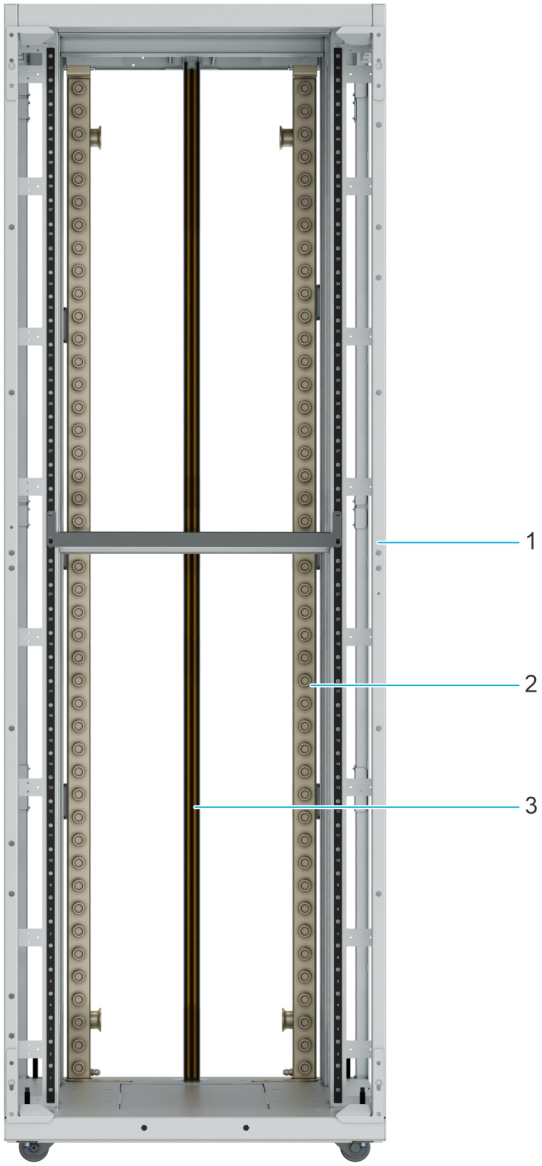

Figure 2. IR7044 Front view

1. Rack frame
2. Blind-mate DLC Manifold
3. Power Busbar

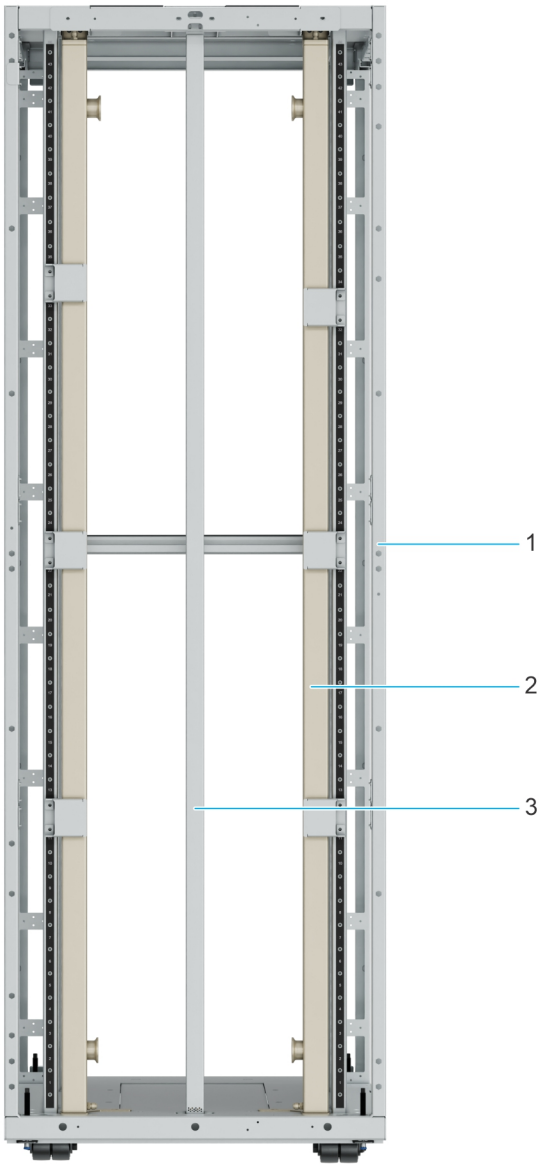

Figure 3. IR7000 rear view

1. Rack frame
2. Blind-mate DLC Manifold
3. Power Busbar

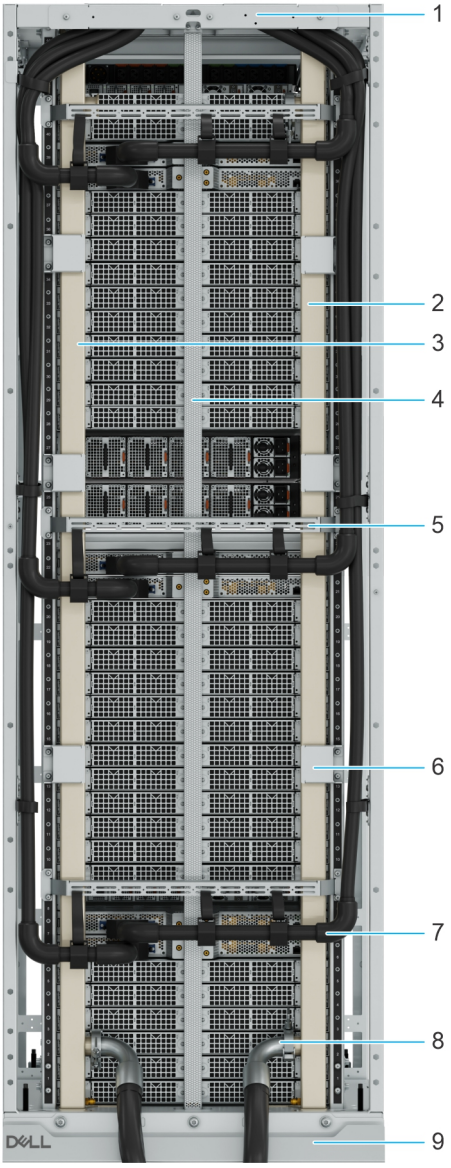

Figure 4. IR7044 rear view with full configuration

1. Rear cable tray bracket
2. Blind-mate DLC Manifold
3. Blind-mate DLC Manifold
4. Power Busbar
5. Strain Relief Bar
6. DLC Manifold bracket
7. Powerwhip
8. DLC Manifold hose
9. Rack Stabilizer

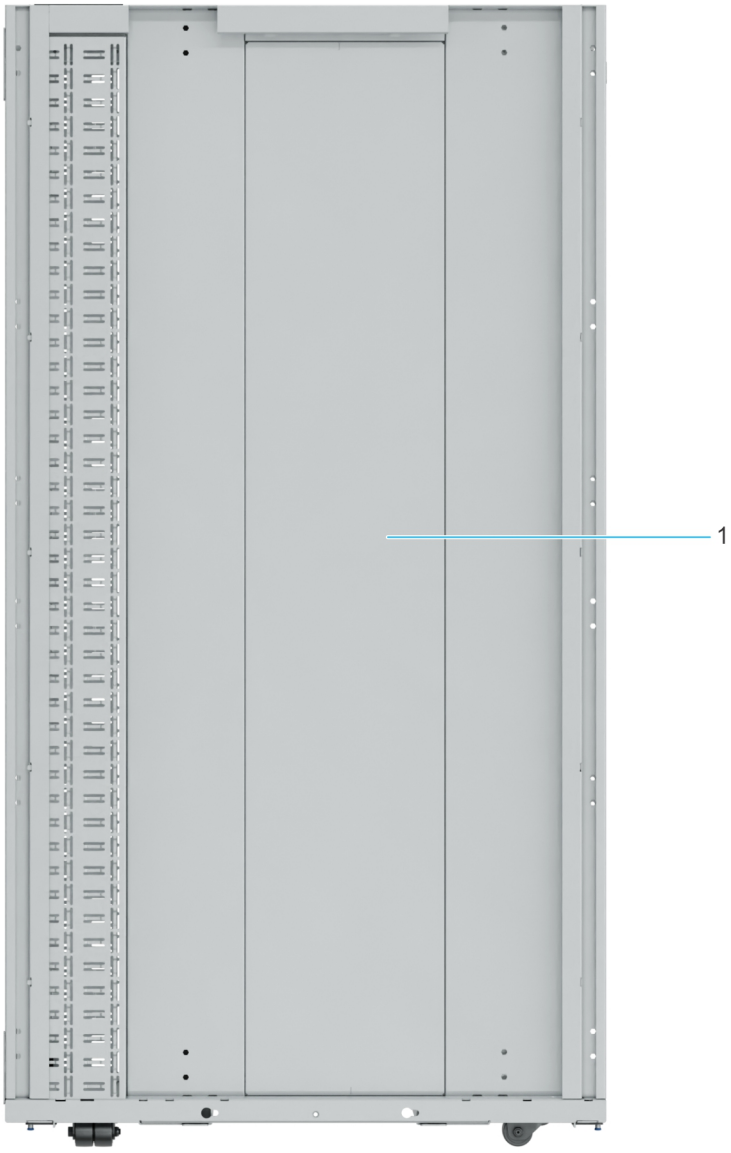

Figure 5. IR7000 side view

1. Side panel

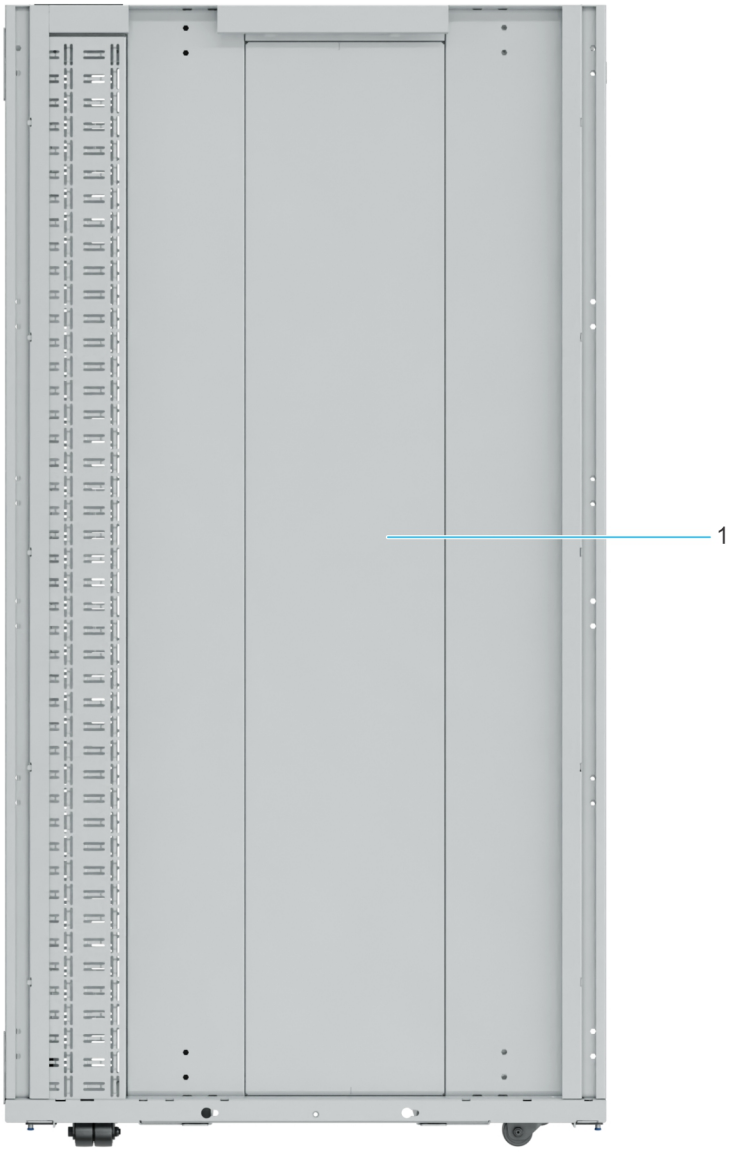

Figure 6. IR7000 side view

1. Side panel

Rack frame dimensions and Weight

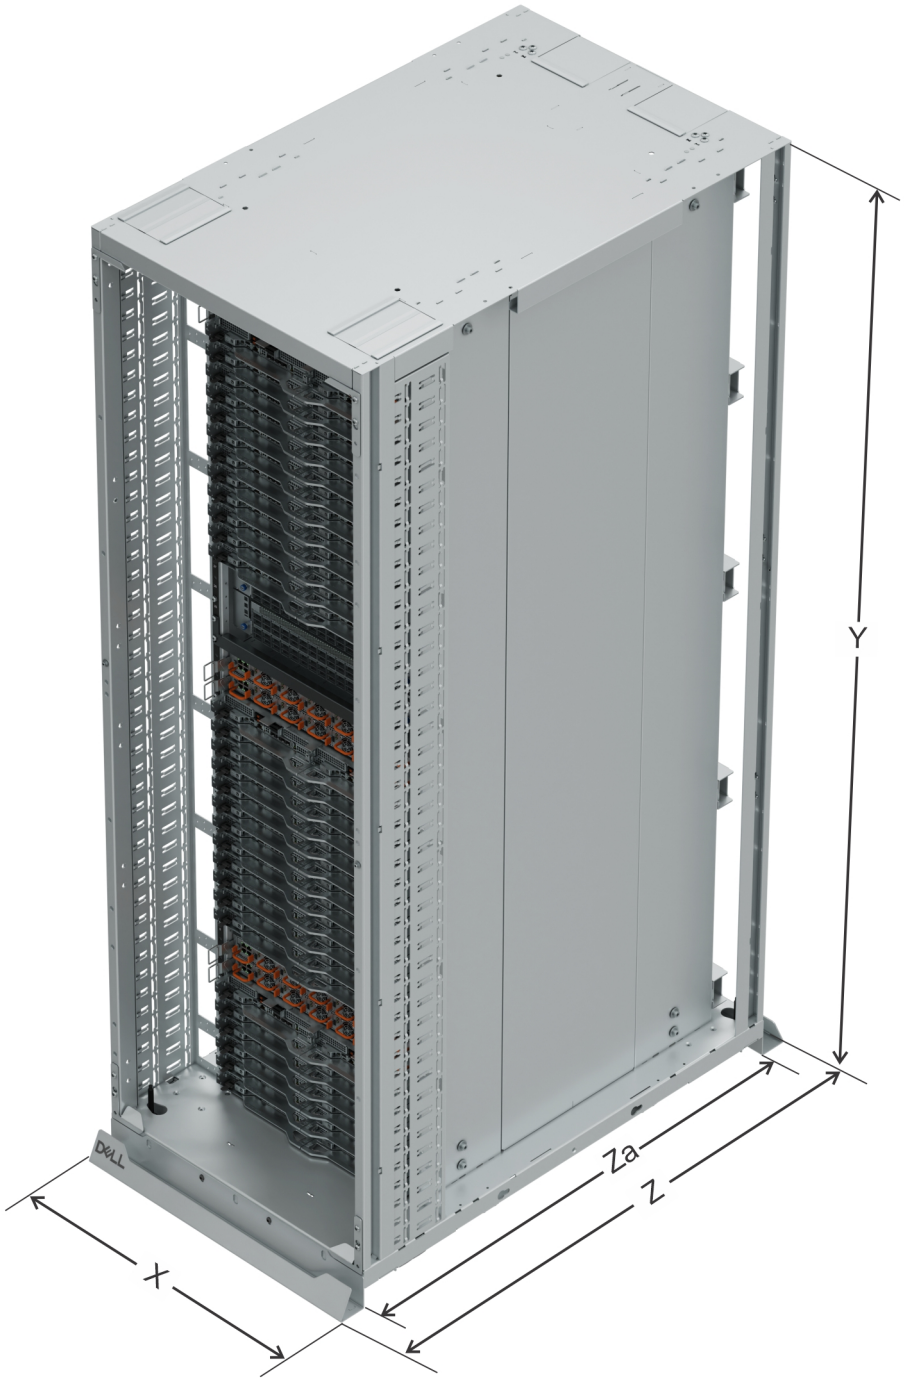

Figure 7. Rack frame dimensions

Table 1. PowerEdge IR7044 and IR7050 Rack dimensions

Type	Rack OU	Height (Y)	Width (X)	Frame depth (Za)	Depth (Z) (with Front & Rear Doors)
Type 1	44	2290 mm (90.15 inch) (Maximum)	750 mm (29.52 inch)	1200 mm (47.24 inch) (Maximum)	1345 mm (52.95 inch) (Maximum)

Table 1. PowerEdge IR7044 and IR7050 Rack dimensions (continued)

Type	Rack OU	Height (Y)	Width (X)	Frame depth (Za)	Depth (Z) (with Front & Rear Doors)
Type 2	50	2578 mm (101.49 inch) (Maximum)	750 mm (29.52 inch)	1200 mm (47.24 inch) (Maximum)	1345 mm (52.95 inch) (Maximum)
Shock Pallet	44 and 50	178 mm (7 inch)	1150 mm (45.27 inch)	1460 mm (57.48 inch)	NA

Table 2. PowerEdge IR7044 and IR7050 Rack weights

Type	Rack OU	Weight
Type 1	44	Max Rack Payload: 1600kg/3530lbs Busbar: 30kg/67lbs Manifold: 55kg/121lbs Usable payload capacity: 1515kg/3342lbs Rack only: 277kg/611lbs
Type 2	50	Max Rack Payload: 1900kg/4189lbs Busbar: 32kg/71lbs Manifold: 64kg/142lbs Usable payload capacity: 1804kg/3976lbs Rack only: 277kg/819lbs
Shock Pallet	44 and 50	70 kgs (154.32 lbs)

Power Busbar overview

The power Busbar distributes DC power from Powershelf to IT Gears mounted in the rack through convenient blind-mate power connections.

- Busbar rating : 1400 A
- Busbar operating voltage : 48 V defined by OCP ORv3

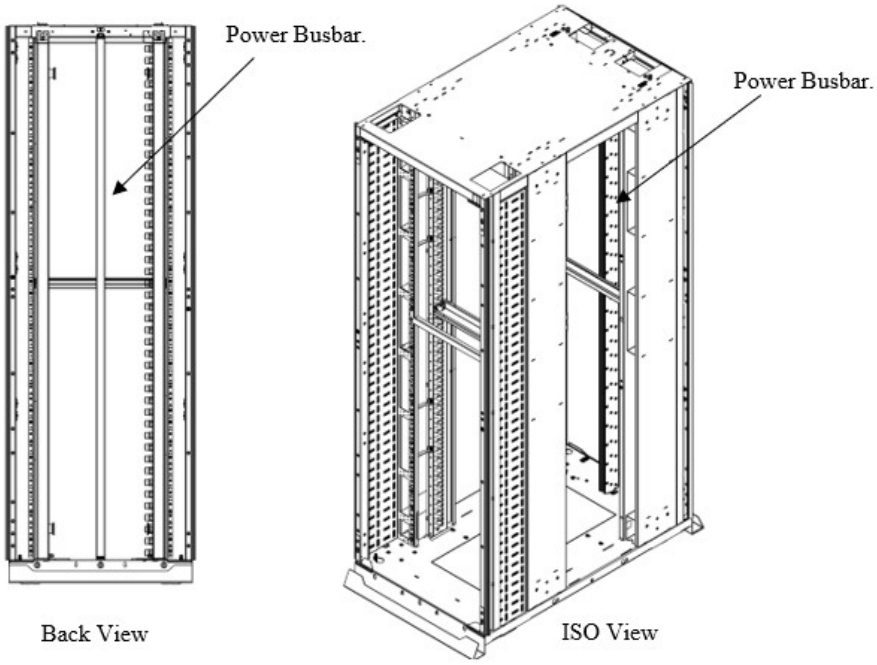

Figure 8. Power Busbar location

Direct Liquid Cooling (DLC) Manifold Overview

DLC Manifold features are:

- 21-inch DLC Manifolds
- Blind-mate Liquid Quick Disconnectors
- Liquid connects at every OU (2 QDs/ OU)
- Utilizes ORv3 spec Manifold mounting interfaces
- Heat capture: Up to 10 kW DLC thermal heat capture per OU

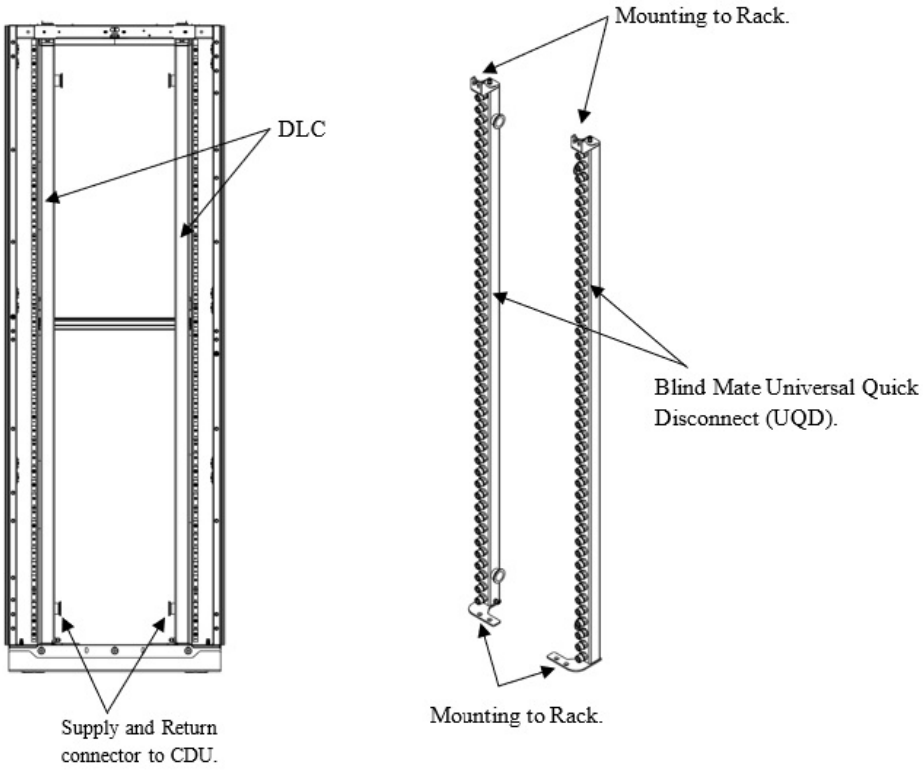

Figure 9. DLC Manifold

Figure 10. DLC mounted on Rack

Figure 11. Direct Liquid Cooling Quick Disconnect (DLC QD) (Left)

Figure 12. Direct Liquid Cooling Quick Disconnect (DLC QD) (Right)

Cooling Distribution Unit (CDU)

NOTE: For more information about the list of supported CDUs, contact your Sales team.

Blind mating for power busbar and Direct Liquid Cooling Universal Quick Disconnect (DLC UQD)

Steps for blind mating the compute enclosure with power busbar and DLC are as follows:

NOTE: Remove the Universal Quick Disconnect (UQD) caps from the system and the Direct Liquid Cooling (DLC) connector in the rack.

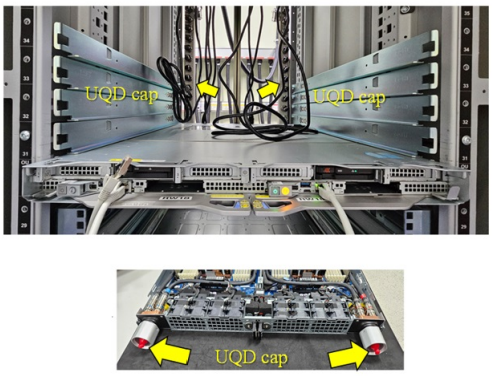

Figure 13. Removing the UQD caps

1. Holding both the levers push the enclosure into the rack.
2. Rotate the levers and push the system until the UQD connectors and power bus bar connector located on the system is engaged with UQD connectors and power bus bar connector rack the rack.
3. Push the compute enclosure holding the handle to engage with DLC manifold and power busbar.

NOTE: It is advised to have two persons to lift the system and placing into the rack due to the system being heavy.

Figure 14. Server installation in the rack

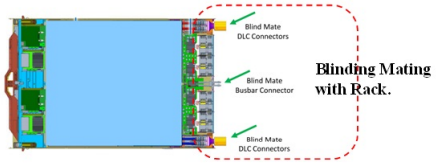

Figure 15. Blind mating for power Busbar and DLC UQD

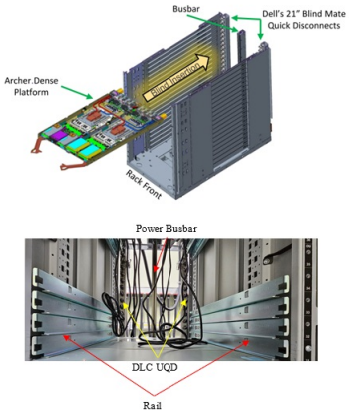

Figure 16. Blind insertion of he compute enclosure

2

Powershelf overview

The M7725 does not support the Power supply units in its enclosure. The power is provided by the six 5500 W power supply units that are installed in the power shelf which provides 51 V DC output to the power busbar which is used by the compute nodes and other components that are installed in the rack.

NOTE: For more information about the Power Management Controller User's Guide, Application Note, and FAQs, contact the Sales team.

Power shelf feature and specifications

- High Efficiency 5.5kW PSU; 97.2% peak / 95.0% @ full load
- Flexible AC input options (Single or Three Phase) with dual input connector
- Hot-pluggable Power Management Controller for monitoring and control
- DC output selects power clip for quick connection to rack power busbar

Table 3. Power shelf input

Configuration	Min	Rated	Max
Star connection (WYE, 3 Ohms 5 W	311 V	346V/480V	528V
Delta connection (Delta, 3 Ohms, 4 W	180V	200V - 277V	305 V
Single phase 1 Ohm	180 V	200V - 277V	305 V
Frequency (Hz)	47	50/60	63
Efficiency	97.2% peak/ 95.0% at full load		
Hold up time	20 ms at 100% load		

Table 4. Power shelf output

Output	
Output power	33 kW
Output voltage (no load/ full load)	49 V / 50 V
Output current	674 A

Figure 17. Power shelf front view

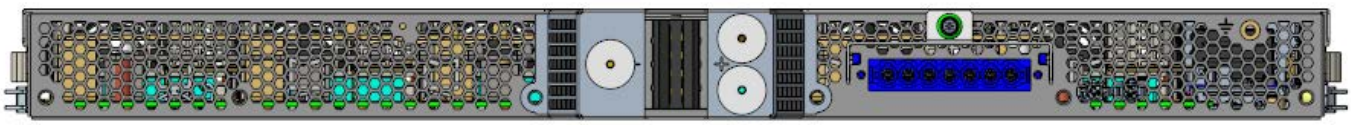

Figure 18. Power shelf rear view

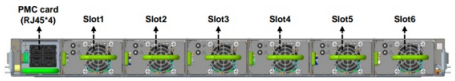

Figure 19. PSU slots

Main input connection

The power shelf has 7 pin 60 A AC input connector

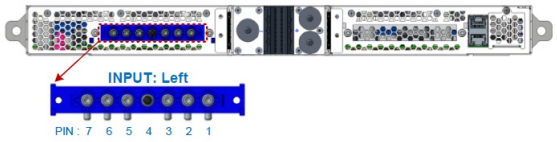

Figure 20. Power shelf input connector

Input fusing

The shelf does not provide any input fusing. Upstream Circuit Breaker with 60A (need to check) per phase recommended. The inserted power supplies are furnished with non-user replaceable 40A fuse(s) in series to line and neutral of power supply input.

Main output connector

The power shelf has a clip connector to mate with the 50 V bus bars which are mounted in rack. Clip has 3 mm floating ability, which can make up for unaligned amount on bus bars.

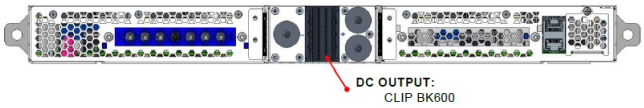

Figure 21. Power shelf output connector

Output grounding

The main output is isolated from the chassis ground.

Hot plug

The Controller is designed to be hot unplugged as well as hot plugged during Shelf operation. The Shelf will not stop operating. In case of a catastrophic fault the Hot plug controller has an internal OCP protection.

Topics:

- Power Management Controller (PMC) module
- Power Shelf dimensions
- Power Shelf weight specifications
- Power shelf environmental specifications
- Power supply unit specifications

Power Management Controller (PMC) module

The PMC is a hot-pluggable controller in the shelf that monitors and facilitates the power supply units through a 10/100/1000MB base Ethernet port and can be connected directly to the rack switch. Redfish protocol is required for the Monitoring and control functions. The supplier is responsible for choosing the communication protocol between PMC to power supply units. The PMC (Power Management Controller) provides Ethernet access to the Power Shelf to communicate with the inserted power supplies. The internal communication between PMC and PSUs is I2C with PMBus protocol. The PMC consists of the following block: The PMC module specifications are listed below:

- Shelf controller for system communication with up to 6 PSUs/1U in the Shelf
- PoE (Power over Ethernet) functions using front RJ45 port.
- Open BMC and Redfish implementation
- Ethernet communication up to 1GB speed
- 512 MB DDR4 RAM
- Dual 64MB Boot Flash
- Internal Debug port
- Temperature sensor
- 1 OU power shelf with 50 V DC bus voltage or 12V standby output

NOTE: The PMC can be Monitored, update the firmware, and Power management through Dell OpenManage Enterprise (OME) software. For more information, see Dell OpenManage Enterprise 4.5 User's Guide or later versions at OpenManage Enterprise User Guide.

NOTE: List of supported Power Shelves devices:

- Delta 21-inch power shelves 33kW 60A
- Lite-On 21-inch power shelves 33kW 60A

Figure 22. Power management controller (PMC)

Reset button and LED

Reset button and LEDs are on the front panel of PMC module. With the reset buttons, it is possible to reset the PMC controller. The LEDs next to the reset button shows the actual status of the PMC.

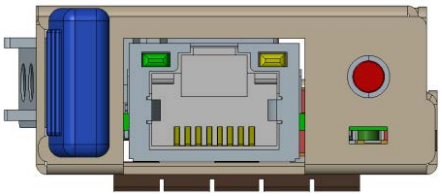

Figure 23. PMC module

LED status

Table 5. LED status

LED	Status
Green	Solid green: Power OK (Voltage is range specification) BMC boot complete
	Blinking green: BMC firmware update in progress
Yellow	Solid yellow: Over temperature warning

Reset button and LEDs are on the front panel of PMC module. With the reset buttons, it is possible to reset the PMC controller. The LEDs next to the reset button shows the actual status of the PMC.

Power Shelf dimensions

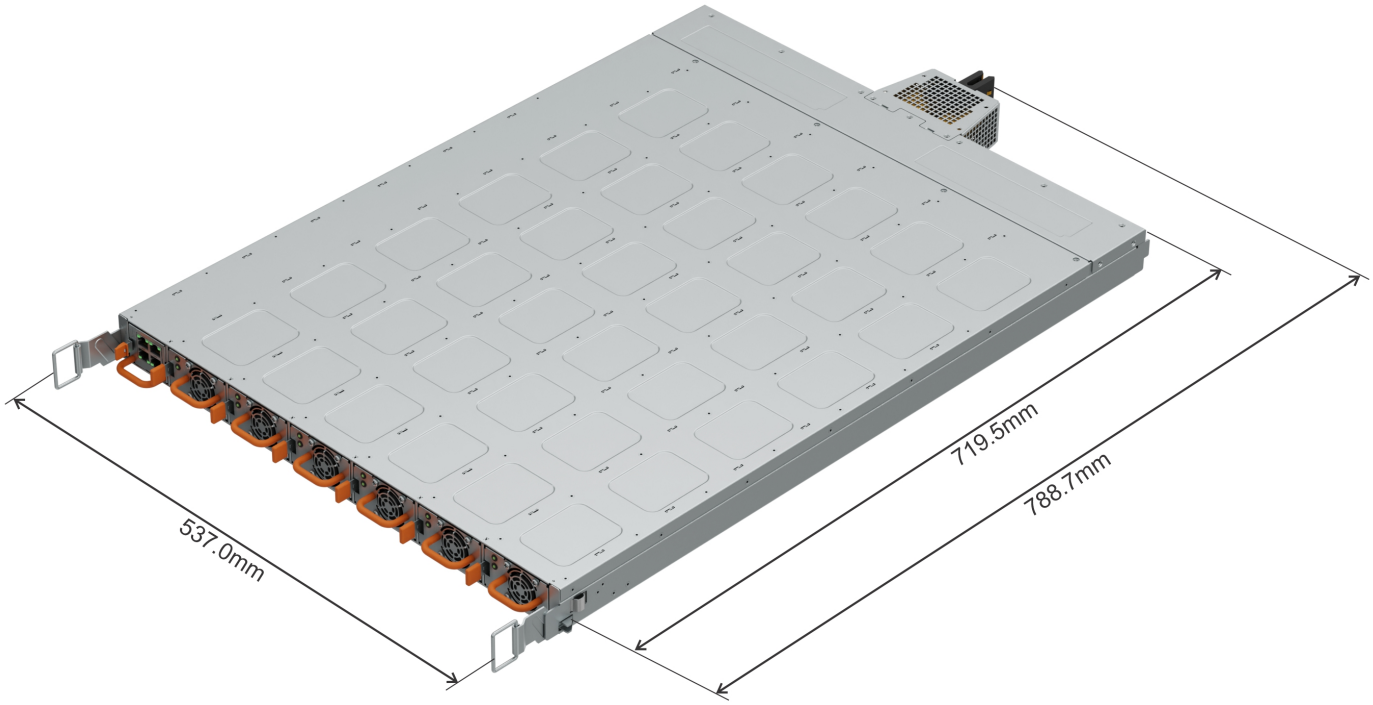

Figure 24. Power Shelf dimensions - 21 inch

Table 6. Power Shelf dimensions 21 inch

Height	Width	Depth	
45.7 mm (1.79 inch)	537 mm (21.14 inch)	719.5 mm (28.32 inch) Shelf latch to shelf rear	788.7 mm (31.05 inch) Shelf latch to busbar connector

Power Shelf dimensions 19 inch

Table 7. Power Shelf dimensions 19 inch

Height	Width	Depth	
43.6 mm (1.71 inch)	495 mm (19.48 inch)	726.5 mm (28.6 inch) Shelf latch to shelf rear	839.5 mm (33.05) Shelf latch to busbar connector

Power Shelf weight specifications

Table 8. Power Shelf weight

Power Shelf	Weight
6 x 5500 W PSUs	38 kg (83.77 lb)

Power shelf environmental specifications

Temperature

- Operational or cold aisle (inlet) temperature: -5°C to +45°C
- Non-operational: -40°C to +85°C

Altitude

- Operational: 0-3050m (10,000ft) above sea-level
- Non-operational : 0-12000m

Thermal margin

Component thermal margin of ≥7% or ≥5°C under worst-case condition (100% load, input voltage lowest, output voltage lowest) and up to 30°C inlet/ambient and 3050m (10,000ft) above sea-level. Target whichever margin value is larger.

Component thermal margin of ≥4% or ≥3°C under worst-case condition (100% load, input voltage lowest, output voltage lowest), at greater than 30°C inlet/ambient and up to 3050m (10,000ft) above sea-level. Target whichever margin value is larger.

Power supply unit specifications

The PSU specifications are as follows:

Table 9. Power supply unit specifications

Key feature	Rating
Input voltage range	200-277 V, 50-60 Hz
Output power	5500 W
Output Voltage (No load/Full load)	50 V/49 V
Output Current	112 A
Efficiency	> 97.5% peak
Operating Temperature	-5°C~45°C
Communication	I2C
Safety Standards	CB/TUV/UL /cUL
EMC	Class A
PSU Dimension	640 mm (L) x 73.4 mm (W) x 40 mm (H)mm

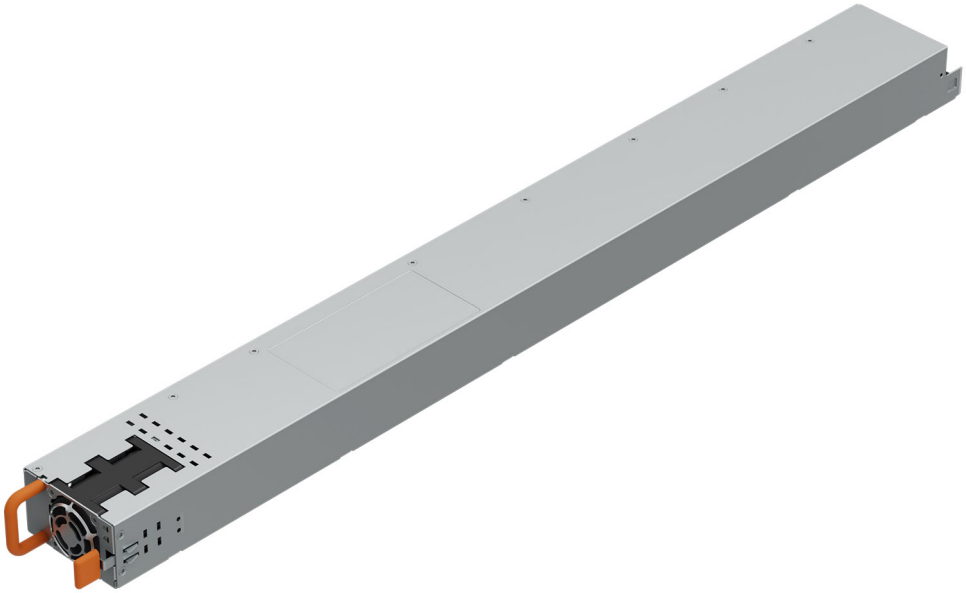

Figure 25. Power supply unit

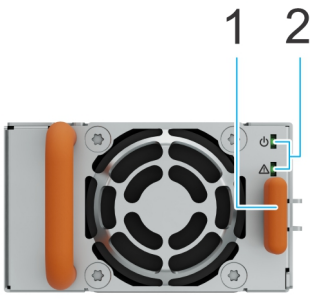

Figure 26. Front view of the PSU

1. PSU handle
2. PSU status LEDs

3

Power whips

A Power Whip is a high-voltage electrical cable assembly used to distribute power from a Power Distribution Unit (PDU) to IT equipment within a data center rack. It plays a crucial role in ensuring efficient and safe power delivery to servers, switches, and other critical components.

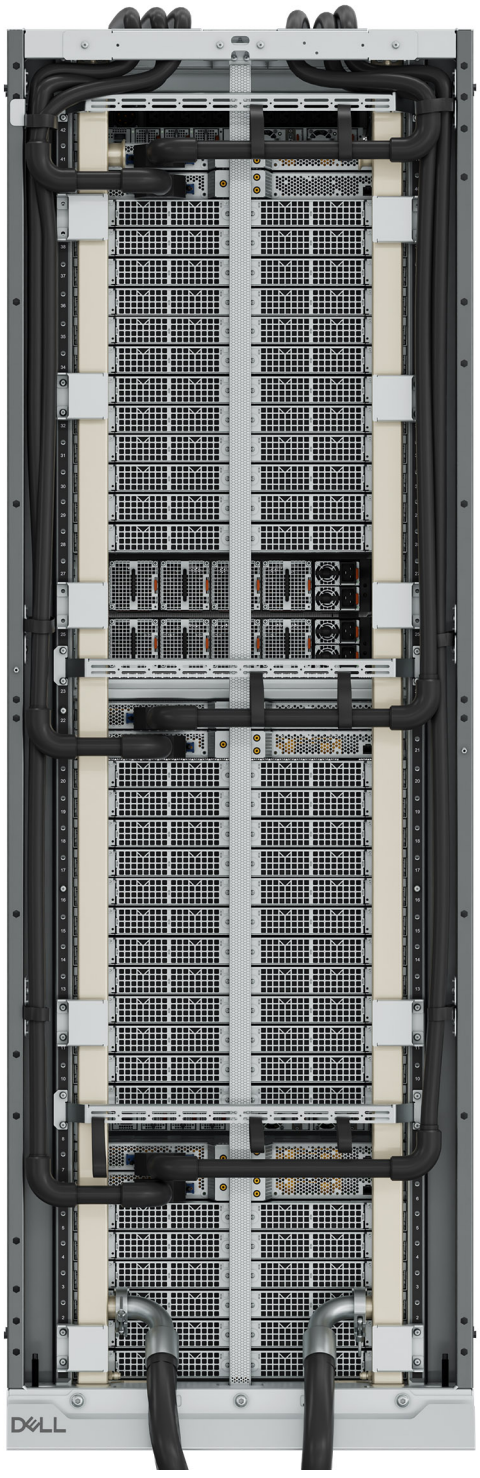

Figure 27. Rear view of IR7000 rack with power whips

Topics:

Key components of power whip

- **Conduit :** Protects the internal wiring from damage and interference.
- **Conductors (Wires) :**Typically made of copper or aluminum to carry electrical current.
- **Connectors :** Custom-fitted to PDUs and power outlets to ensure secure connections.
- **Jacket :** Insulated covering for safety and durability.

Types of power whips

- **Flexible Metal Conduit (FMC) :** Provides flexibility in cable routing.
- **Liquid-Tight Flexible Metal Conduit (LFMC) :** Offers moisture resistance for damp environments.
- **Rigid Metal Conduit (RMC) :** Used for high-security installations requiring extra protection.

Power whip specifications

- **Voltage Ratings :** Commonly available in 120V, 208V, 240V, or higher for data centers.
- **Ampere Ratings :** Typically ranges from 15A to 60A, depending on power requirements.
- **Phase Configuration :** Can be single-phase or three-phase for higher power efficiency.
- **Length :** Varies based on rack layout and distance from the PDU.

Benefits of power whips in IT racks

- **Reliable Power Distribution :** Ensures consistent power delivery to critical IT equipment.
- **Reduced Cable Clutter :** Helps maintain an organized rack setup.
- **Customization :** Available in different configurations to match PDU and equipment needs.
- **Safety Compliance :** Designed to meet electrical codes and standards for safe operation.

Installation & best practices

- Ensure power whips are rated for the rack's voltage and amperage requirements.
- Properly secure conduits to avoid strain on connections.
- Maintain adequate airflow around cables to prevent overheating.
- **NOTE:** Regularly inspect for wear and replace damaged power whips to prevent failures.

4

Management switch overview

The IR7000 is equipped with management switch for networking within the rack with DELL S5248-ON management switch

Table 10. Management switch dimensions

Height	Width	Depth
43.6 mm (1.71 inch)	434 mm (17.08 inch)	460 mm (18.11 inch)

Figure 28. Management switch

NOTE: Installing the switch in the IR7000 rack requires a 19-inch to 21-inch rail conversion kit.

5

Fabric switch overview

The IR7000 supports the following fabric switches:

- Dell Z9664F-ON 64 x 400 Gb
- Dell Z9864F-ON 64 x 800 Gb
- NVIDIA Quantum2 QM9700- NS2R

Table 11. DELL Z9664F-ON 64 x 400 Gb switch dimensions

Height	Width	Depth
87.5 mm (3.44 inch)	438.5 mm (17.26 inch)	630 mm (24.80 inch)

Figure 29. Dell Z9664F-ON 64 x 400 Gb fabric switch

Table 12. Dell Z9864F-ON 64 x 800 Gb switch dimensions

Height	Width	Depth
87.3 mm (3.43 inch)	438.5 mm (17.26 inch)	630 mm (24.80 inch)

Figure 30. Dell Z9864F-ON 64 x 800 Gb fabric switch

Table 13. NVIDIA Quantum2 QM9700- NS2R switch dimensions

Height	Width	Depth
43.6 mm (1.71 inch)	438 mm (17.24 inch)	660 mm (25.98 inch)

Figure 31. NVIDIA Quantum2 QM9700- NS2R fabric switch

NOTE: Installing all the switches in the IR7000 rack requires a 19-inch to 21-inch conversion kit.

6

Rack Power Distribution Unit (rPDU)

The IR7000 is supports Vertiv VP4UU1A0 power distribution unit to provide power to the components installed in the rack.

Table 14. rPDU dimensions

Height	Width	Depth
44 mm (1.73 inch)	432 mm (17 inch)	229 mm (9.01 inch)

Figure 32. rPDU front view

Figure 33. rPDU rear view

The rPDU has a universal inlet and plug form (3P+N+E) (IP44).

7

PowerEdge M7725 system configurations and features

The PowerEdge M7725 system is a two node in a 1 OU (Open rack unit) installed in M7701 chassis that supports:

- Two 5 th Generation AMD EPYC 9005 Series processor with up to 128 Zen5 cores and 192 Zen5c cores per processor per node
- 24 DDR5 DIMM slots per node
- Up to 2 x E3.S NVMe Gen5 (SSD front) drives per node

NOTE: 1 OU = 48 mm (1.88 inch) height, and 538.98 mm (21.22 inch) width.

NOTE: M7701 1 OU chassis supports 2 x M7725 nodes.

NOTE: For more information about how to hot swap NVMe PCIe SSD device, see the *Dell Express Flash NVMe PCIe SSD User's Guide* at Dell Support page **> Browse all products** > **Infrastructure** > **Data Center Infrastructure** > **Storage Adapters & Controllers** > **Dell PowerEdge Express Flash NVMe PCIe SSD** > **Select This Product** > **Documentation** > **Manuals and Documents**.

NOTE: The speed of E3.S drives is also dependent on the storage controller capability and not just the drive speed.

CAUTION: Do not install GPUs, network cards, or other PCIe devices on your system that are not validated and tested by Dell. Damage caused by unauthorized and invalidated hardware installation will null and void the system warranty.

Topics:

Key workloads

The Dell PowerEdge M7725is purpose-built for dense compute. The key workloads for M7725 are :

- Genomics Sequencing and Analysis
- Computational Fluid Dynamics Simulations
- Computer-Aided Engineering
- Weather Forecasting Models
- Financial Risk Modeling
- High-Frequency Trading
- Digital Visual Effects and Animation

New technologies

The PowerEdge M7725 is capable of handling demanding workloads and applications, such as high-performance computing (HPC), high-frequency trading, digital visual effects and animation.

Table 15. New technologies

Technology	Detailed Description
AMD EPYC 9005 Series processor	Core count: Up to 128 Zen5 cores and 192 Zen5c cores per processor per node

Table 15. New technologies (continued)

Technology	Detailed Description
	4 nm processor technology
	AMD Inter-chip global memory interconnect (xGMI) up to 4 G-links with speeds up to 32Gbps
	Speeds up to 3.6 GHz
	Maximum TDP: 500 W
6400 MT/s DDR5 Memory	Up to 12 channels with 1 DPC per CPU and 24 DIMMs in total
	Supports RDIMM, 3DS DIMM DDR5 with ECC up to 6400 MT/s
PCIe Gen	Gen5 @32 GT/s
PCIe Slot	Up to two PCIe slots with x16 lanes per node
Flex I/O	Front I/O with DC-SCM card : ● 1 x Dedicated iDRAC ethernet port ● 1 x USB 3.0 Type-A port ● 1 x USB 2.0 Type-C port ● 1 x mini display port
CPLD 1-wire	Support payload data of Riser, BOSS N1 Modular DC MHS, BP, and Rear IO to BIOS and iDRAC

8

Chassis views and features

Topics:

- Front view of the system
- Rear view of the system
- System configurations - inside view for PowerEdge M7725 system
- Quick Resource Locator

Front view of the system

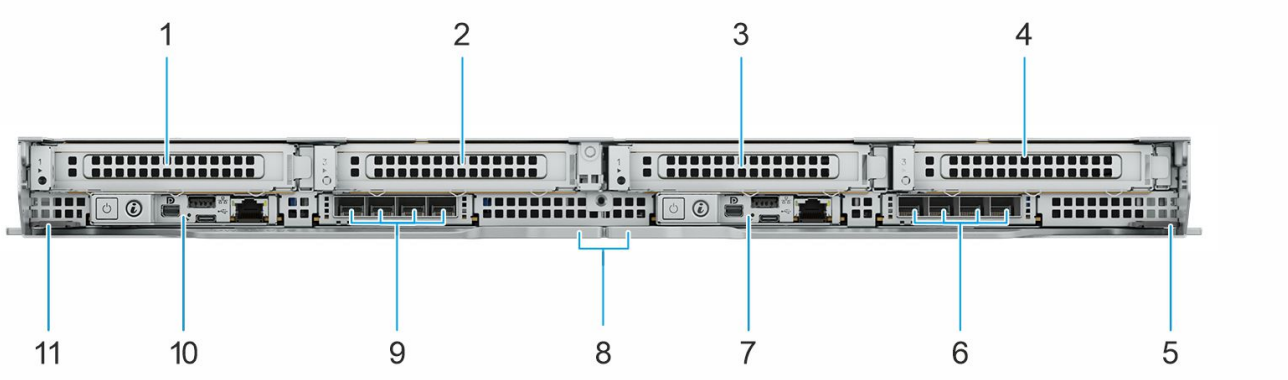

Figure 34. Front view with DC-SCM + OCP 3 NIC + 2 x Full height risers per node

1. Riser F1B (Full height, node 1)
2. Riser F2B (Full height, node 1)
3. Riser F1B (Full height, node 2)
4. Riser F2B (Full height, node 2)
5. Information tag (node 2)
6. OCP NIC (node 2)
7. DC-SCM module (node 2)
8. Chassis lever
9. OCP NIC (node 1)
10. DC-SCM module (node 1)
11. Information tag (node 1)

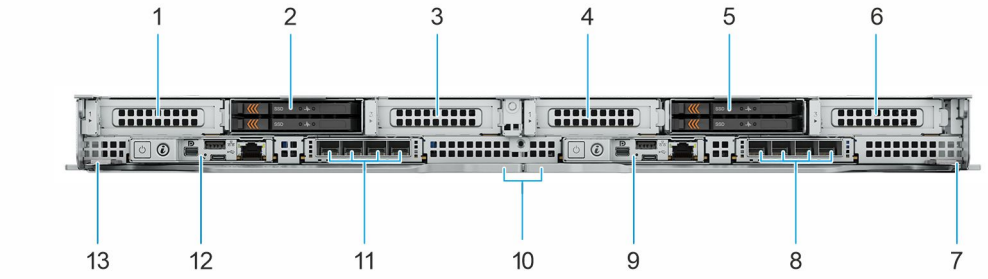

Figure 35. Front view with DC-SCM + OCP 3 NIC + 2 x Low profile risers + 2 x E3.S drives per node

1. Riser F1A (Low profile, node 1)

2. E3.S drive module (node1)
3. Riser F2A (Low profile, node1)
4. Riser F1A (Low profile, node 2)
5. E3.S drive module (node2)
6. Riser F2A (Low profile, node2)
7. Information Tag (node 2)
8. OCP NIC (node 2)
9. DC-SCM module (node 2)
10. Chassis lever
11. OCP NIC (node 1)
12. DC-SCM (node 1)
13. Information Tag (node 1)

Rear view of the system

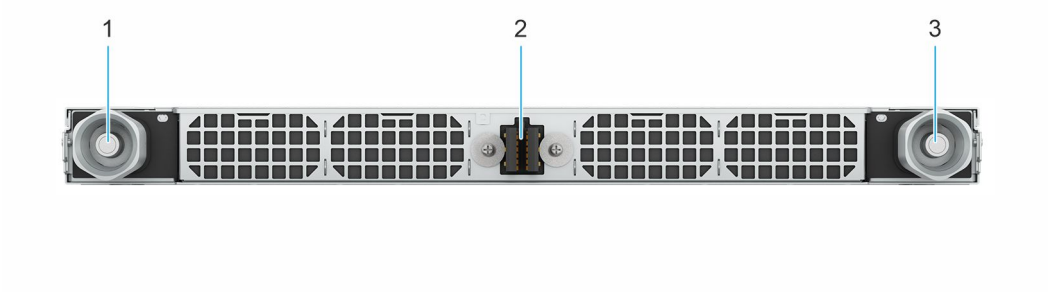

Figure 36. Rear view of the system showing blind mate connections to rack power and liquid cooling interfaces and 4 x 40 mm fans per node

Table 16. Rear view of the system

Item	Ports, panels, or slots	Icon	Description
1	IMM return outlet	N/A	Enables the coolant liquid to flow out of the system to the DLC manifold located on the rack.
2	Power bus bar connector	N/A	Enables to get power supply from the power bus bar located on the rack.
3	IMM supply inlet	N/A	Enables the coolant liquid to flow into the system from the DLC manifold located on the rack.

System configurations - inside view for PowerEdge M7725 system

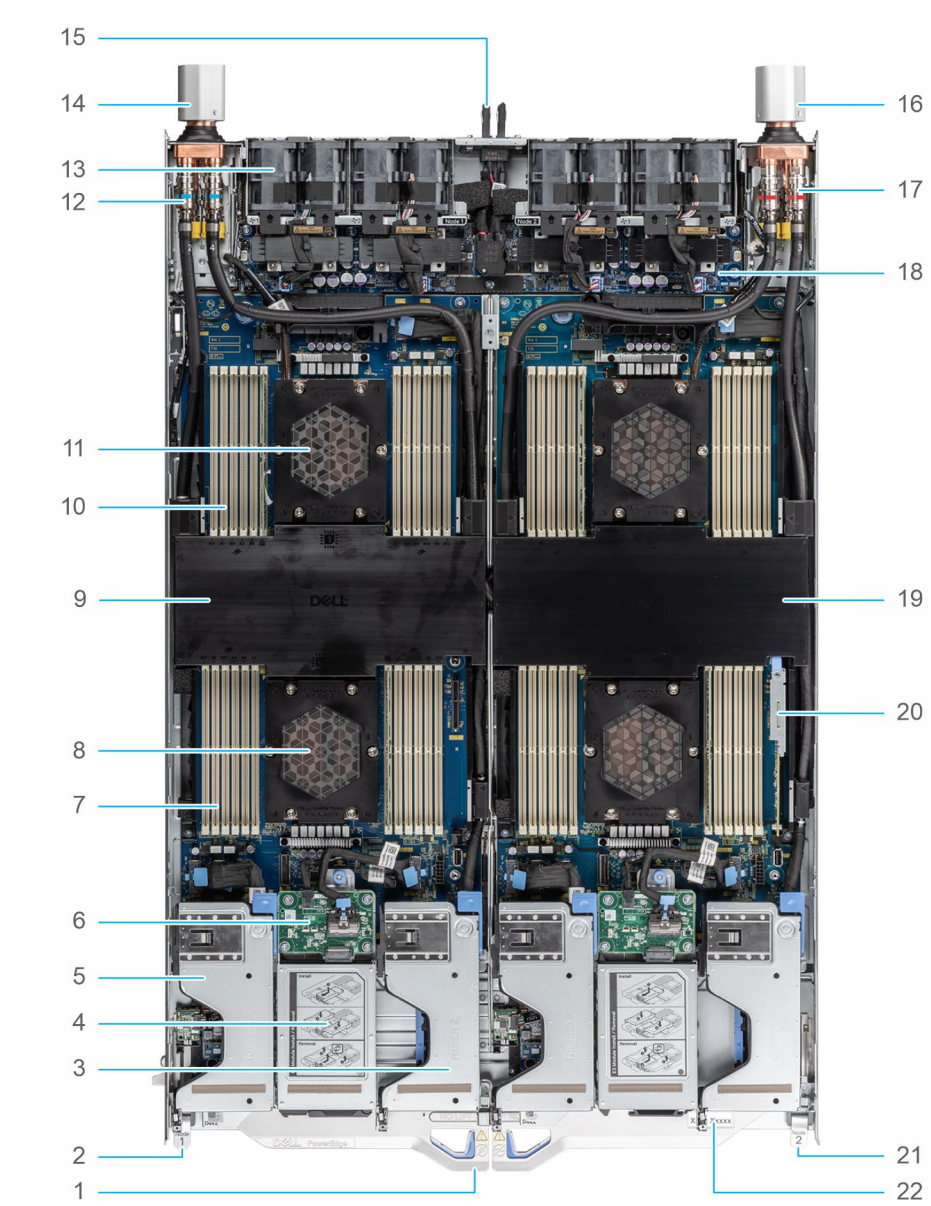

Figure 37. Inside view PowerEdge M7725 system with low profile risers and E3.S drive modules

1. Enclosure lever
2. Node 1 information tag
3. RIser F2A (Low profile)
4. E3.S SSD module
5. RIser F1A (Low profile)
6. E3.S SSD backplane
7. Memory modules (CPU 1)
8. Direct Liquid Cooling (DLC) cold plate for CPU 1
9. Air shroud (Node 1)
10. Memory modules (CPU 2)
11. Direct Liquid Cooling (DLC) cold plate for CPU 2
12. IMM (Internal Manifold Module) coolant inlet
13. Cooling fans
14. Universal Quick Disconnect Blind (UQDB) inlet
15. Power bus bar connector
16. Universal Quick Disconnect Blind (UQDB) outlet
17. IMM (Internal Manifold Module) coolant outlet
18. Power Distribution Board (PDB)
19. Air shroud (Node 2)
20. BOSS-N1 Modular DC-MHS
21. Node 2 information tag
22. M7701 chassis service tag

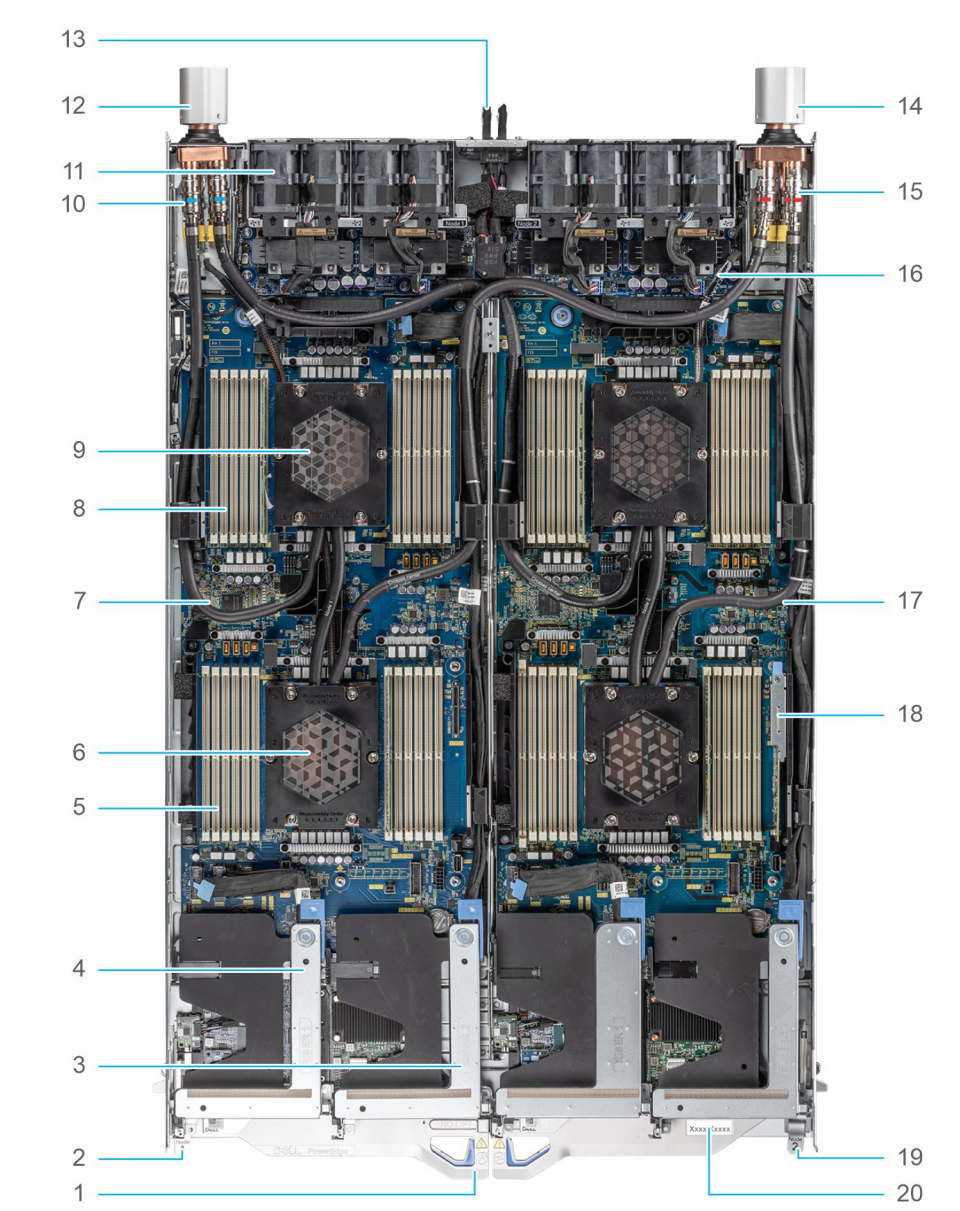

Figure 38. Inside view of M7725 with full height risers

1. Enclosure lever
2. Node 1 information tag
3. Riser F2B (Full Height)
4. Riser F1B (Full Height)
5. Memory modules (CPU 1)
6. Direct Liquid Cooling (DLC) cold plate for CPU 1
7. Direct Liquid Cooling (DLC) tube
8. Memory modules (CPU 2)
9. Direct Liquid Cooling (DLC) cold plate for CPU 2
10. IMM (Internal Manifold Module) coolant inlet
11. Cooling fans
12. Universal Quick Disconnect Blind (UQDB) inlet
13. Power bus bar connector
14. Universal Quick Disconnect Blind (UQDB) outlet
15. IMM (Internal Manifold Module) coolant outlet
16. Power Distribution Board (PDB)

17. Direct Liquid Cooling (DLC) tube
18. BOSS-N1 DC-MHS module
19. Node 2 information tag
20. M7701 chassis service tag

Quick Resource Locator

The QR code on everything (SILs, GSG, Owner's Manual except on the EST) is a generic QR code for M7725 that leads to a webpage for that product. That webpage has links for things like setup and service videos, iDRAC manual, and other things that apply to the platform. The QR code on the EST is unique and specific to that service tag and will contain the Service Tag number and the iDRAC password. The label and the QR code within it are printed on demand at the L10 factories. This QR code links to a webpage that shows the exact configuration as built for that customer, and the specific warranty purchased. It is one click away from the same content of generic information that applies to M7725 that is available in the other QR codes.

Figure 39. Quick Resource Locator for PowerEdge M7725 system

9

Chassis dimensions

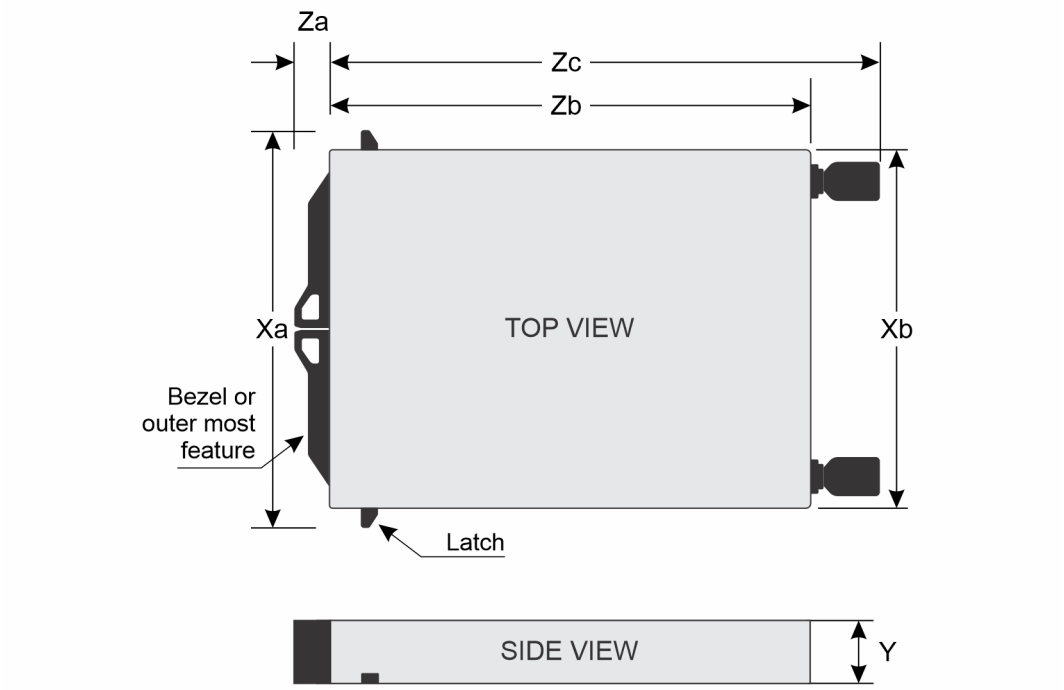

Figure 40. Chassis dimensions

Table 17. PowerEdge M7725 chassis dimensions

Xa	Xb	Y	Za	Zb	Zc
560 mm (22.05 inches) **NOTE:** The front bezel is not supported and no rack ears	537.0 mm (21.14 inches)	46.45 mm (1.83 inches)	71.81 mm (2.83 inch) **NOTE:** The front bezel is not supported	789.85 mm (31.10 inches) Hook to chassis rear	848.80 mm (33.41 inches) Hook to IMM module alighment cup

10

System weight

Table 18. PowerEdge M7725 system weight

System configuration	Maximum weight (with all drives/SSDs)
4 x16 Full height Riser configuration	26.37 kg (58.13 pounds)
4 x8 Low profile + 4 x E3.S configuration	27.9 kg (61.5 pounds)

11

Processor

Topics:

- Processor features

Processor features

The AMD 5th Generation AMD EPYC 9005 Series processor is next generation data center CPU supporting socket compatibility with EPYC™ 9004 “Genoa” in SP5+ socket infrastructure. Based on of AMD’s new enhanced “Zen5” & “Zen5C” CPU core with integrated I/O controllers, AMD EPYC™ “Turin” SOC offers significant performance improvement from current generation production and the best performance per price and lowers TCO through an optimal balance of compute, memory,I/O, andd security.

The following lists the features and functions that are in the I5th Generation AMD EPYC 9005 Series processor offering:

Compute "Zen5" cores

- Up to 128 Zen5 cores / 2 × threads per socket
- Up to 32 MB L3 shared by 8 cores/CCD
- 1 MB L2/core, 32/48 KB instruction/data L1 per core

Compute "Zen5c" cores

- Up to 192 Zen5c cores / 2 × threads per socket
- Up to 32 MB L3 shared by 16 cores/CCD
- 1 MB L2/core, 32/48 KB instruction/data L1 per core

Memory

- 12 DDR5 memory channels up to 6400 MT/s (1DPC)
- RDIMM, 3DS RDIMM
- Dynamic PPR for non-Chipkill DIMMs

Integrated I/O

- PCIe5 supports, peak xGMI3 product speeds up to 32Gps
- Up to 128 lanes of High Speed I/O
- Server Controller Hub (USB, UART, SPI, LPC, I2C, etc.)

Supported processors

The following table shows the 5th Generation AMD EPYC 9005 Series processor SKUs that are supported on the M7725.

Table 19. AMD EPYC 9005 Series processors supported in M7725

Processor	Base Clock Speed (GHz)	Max Clock Speed (GHz)	Cache (M)	Cores	Threads	Turbo	Memory Speed (MT/s)	Memory Capacity	TDP
AMD EPYC ™9475F	3.65	4.8	256	48	96	Turbo	6400	6 TB	400 W
AMD EPYC ™9575F	3.3	5.0	256	64	128	Turbo	6400	6 TB	400 W
AMD EPYC ™9825	2.2	3.7	384	144	288	Turbo	6400	6 TB	390 W

Table 19. AMD EPYC 9005 Series processors supported in M7725 (continued)

Processor	Base Clock Speed (GHz)	Max Clock Speed (GHz)	Cache (M)	Cores	Threads	Turbo	Memory Speed (MT/s)	Memory Capacity	TDP
AMD EPYC ™9845	2.1	3.7	320	160	320	Turbo	6400	6 TB	390 W
AMD EPYC ™9655	2.6	4.5	384	96	192	Turbo	6400	6 TB	400 W
AMD EPYC ™9755	2.7	4.1	512	128	256	Turbo	6400	6 TB	500 W
AMD EPYC ™9965	2 .25	3.7	384	192	384	Turbo	6400	6 TB	500 W

12

Memory subsystem

Topics:

Supported memory

Table 20. Memory technology comparison

Feature	PowerEdge M7725 (DDR5)
DIMM type	RDIMM
Transfer speed	6400 MT/s(1DPC) **NOTE:** Maximum DIMM transfer speed support dependent on CPU SKU and DIMM population
Voltage	1.1 V

Table 21. Supported DIMMs

DIMM PN	Rated DIMM Speed (MT/s)	DIMM Type	DIMM Capacity (GB)	Ranks per DIMM	Data Width	DIMM Volts (V)
9C0R6	6400	RDIMM	128	2	x4	1.1
G9PYX	6400	RDIMM	32	2	x8	1.1
N66RP	6400	RDIMM	64	2	x4	1.1
JRGVT	6400	RDIMM	96	2	x4	1.1

Table 22. Supported memory matrix

DIMM type	Rank	Capacity	DIMM rated voltage and speed	Operating Speed
				1 DIMM per channel (DPC)
RDIMM DDR5 ECC	2 R	32 GB, 64 GB, 96 GB, 128 GB	DDR5 (1.1 V), 6400 MT/s DDR5 (1.1 V), 6400 MT/s	Up to 6400 MT/s

NOTE: The processor may reduce the performance of the rated DIMM speed.

General memory module installation guidelines

To ensure optimal performance of your system, observe the following general guidelines when configuring your system memory. If your system's memory configuration fails to observe these guidelines, your system might not boot, stop responding during memory configuration, or operate with reduced memory.

The memory bus may operate at speeds of 6400 MT/s or lower speed depending on the following factors:

- System profile selected (for example, Performance, Performance Per Watt Optimized (OS), or Custom [can be run at high speed or lower])
- Maximum supported DIMM speed of the processors
- Only DDR5- 6400 MT/s RDIMMs are supported

NOTE: MT/s indicates DIMM speed in MegaTransfers per second.

NOTE: Fault Resilient Memory and memory mirroring are not supported.

The following are the recommended guidelines for installing memory modules:

- All DIMMs must be DDR5.
- The M7725 system supports 1, 2, 4, 6, 8, 10, 12 DIMMs per CPU.
- Do not mix DIMM module types within a memory channel. All must be RDIMM type, with same ECC configuration.
- Do not mix x4 and x8 DIMMs within a memory channel
- Mixing DIMMs of different rank types will change or disable the chip select interleaving feature and may have a performance impact.
- DIMM mixing configurations are not supported. All DIMM slots must be populated with the exact same DIMMs.
- Populate memory module sockets only if a processor is installed.
  - For dual-processor systems, sockets A1 to A12 and sockets B1 to B12 are available.
- In **Optimizer Mode**, the DRAM controllers operate independently in the 64-bit mode and provide optimized memory performance.

Table 23. Memory population rules

Processor	Memory population	Memory population information
Dual processor	A{1}, B{1}, A{2}, B{2}, A{3},B{3}, A{4}, B{4}, A{5}, B{5},A{6}, B{6}, A{7}, B{7} A{8},B{8}, A{9}, B{9}, A{10},B{10}, A{11}, B{11}, A{12}, B{12}	24 DIMMs are supported per node. 1, 2, 4, 6, 8, 10, 12 DIMMs per CPU

- Unbalanced or odd memory configurations result in a performance loss, and the system may not identify the memory modules being installed. Always populate memory channels identically with equal DIMMs for the best performance.

System memory guidelines

The PowerEdge M7725 system supports DDR5 registered DIMMs (RDIMMs). System memory holds the instructions that are started by the processor.

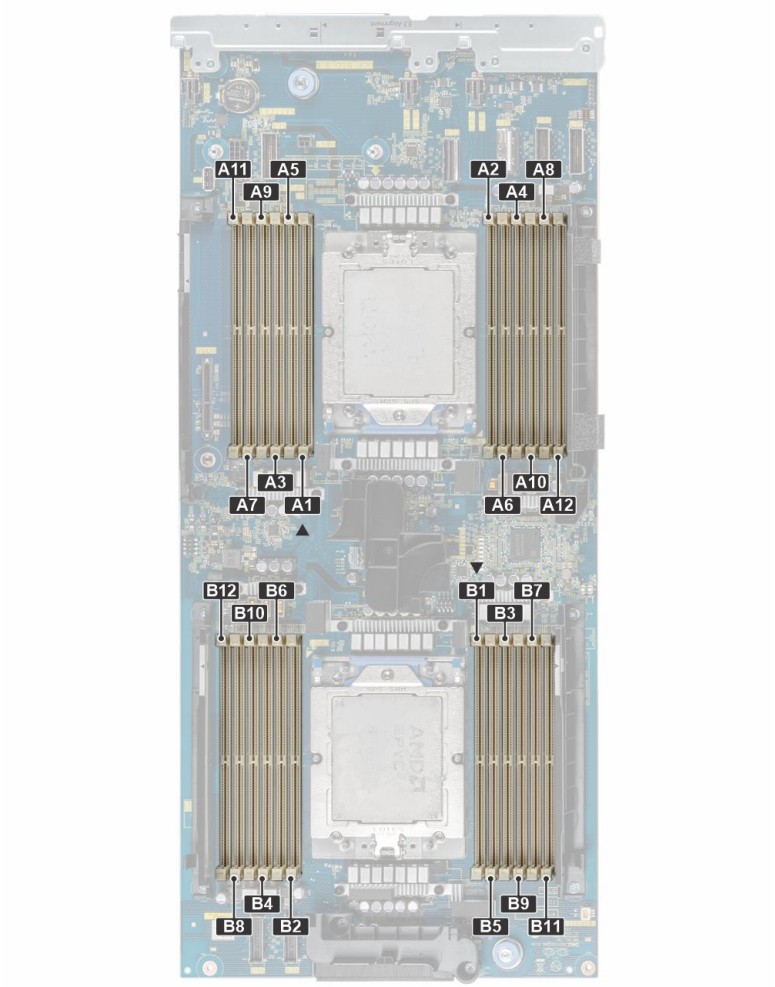

Figure 41. Memory channels

Memory channels are organized as follows:

Table 24. Memory channels

Processor	Channel A	Channel B	Channel C	Channel D	Channel E	Channel F	Channel G	Channel H	Channel I	Channel J	Channel K	Channel L
Processor 1	Slot A1	Slot A5	Slot A3	Slot A9	Slot A7	Slot A11	Slot A2	Slots A6	Slot A4	Slot A10	Slot A8	Slot A12
Processor 2	Slot B1	Slot B5	Slot B3	Slot B9	Slot B7	Slot B11	Slot B2	Slots B6	Slot B4	Slot B10	Slot B8	Slot B12

NOTE: The processor may reduce the performance of the rated DIMM speed.

NOTE: Maximum DIMM transfer speed support dependent on CPU SKU and DIMM population.

Memory RAS features

Reliability, Availability, and Serviceability (RAS) features help keep the system online and operational without significant impact to performance and can decrease data loss and failing due to errors. RAS aids in rapid, accurate diagnosis of faults which require service.

The table below describes the memory RAS features supported on the platform.

Table 25. Supported RAS features

Feature	Description
Demand Scrubbing	Demand scrubbing is the ability to write corrected data back to the memory, once a correctable error is detected on a read transaction. This allows for correction of data in memory at the time of detection, and decreases the chances of a second error on the same address accumulating and causing a multi-bit error condition.
Patrol Scrubbing	Patrol scrubbing proactively searches the system memory repairing correctable errors preventing accumulation of single-bit errors and turning it into an uncorrected error. Patrol scrubbing is accomplished using an engine that generates requests to memory addresses in a stride. The engine will generate a memory request at the Pre-programmed frequency, and the demand scrubbing flow corrects the error, if any. Patrol scrubbing finds opportunities on idle cycles to scrub the memory and get rid of any detectable correctable errors. Patrol scrubs are intended to ensure that data with a correctable error does not remain in DRAM long enough to stand a significant chance of further corruption to an uncorrectable error due to high energy particle error. The IMC will issue a Patrol Scrub at a rate sufficient to write every line once a day. For a maximum channel capacity of 192 GB, this would be one scrub every 26.8 micro-sec. The Patrol Scrub rate is configurable using 16b scrub interval timer.
Permanent Fault Detection (PFD)	PFD is starting with Sapphire Rapids processor. The logic determines if a given fault from DIMM is confined to a single device (Correctable), multi devices (Uncorrectable), or if the fault was transient. The ECC logic makes use of this information to correct the error from a faulty DRAM device.

For RAS modes that require matching DIMM populations, the same slot positions across channels must hold the same DIMM type regarding size and organization. DIMM timings do not have to match, but timings are set to support all DIMMs populated (that is, DIMMs with slower maximum timings force faster DIMMs to the slower of the maximum timing modes).

13

Storage

Topics:

Supported Drives

The table that is shown below lists the internal drives that are supported in M7725.

Table 26. Supported drives

Manufacturer	Form Factor	Interface	Sector size	Drive type	Capacities
Micron	E3.S	PCIe Gen5 NVMe	512e	SSD	3.84 TB
Micron	E3.S	PCIe Gen5 NVMe	512e	SSD	7.68 TB
Micron	E3.S	PCIe Gen5 NVMe	512e	SSD	15.36 TB
Micron	E3.S	PCIe Gen5 NVMe	512e	SSD	3.2 TB
Micron	E3.S	PCIe Gen5 NVMe	512e	SSD	6.4 TB
Kioxia	E3.S	PCIe Gen5 NVMe	512e	SSD	1.92 TB
SK Hynix	E3.S	PCIe Gen5 NVMe	512e	SSD	1.92 TB
Kioxia	E3.S	PCIe Gen5 NVMe	512e	SSD	3.84 TB
Samsung	E3.S	PCIe Gen5 NVMe	512e	SSD	3.84 TB
SK Hynix	E3.S	PCIe Gen5 NVMe	512e	SSD	3.84 TB

Table 26. Supported drives (continued)

Manufacturer	Form Factor	Interface	Sector size	Drive type	Capacities
Kioxia	E3.S	PCIe Gen5 NVMe	512e	SSD	3.84 TB
SK Hynix	E3.S	PCIe Gen5 NVMe	512e	SSD	3.84 TB
Kioxia	E3.S	PCIe Gen5 NVMe	512e	SSD	7.68 TB
Samsung	E3.S	PCIe Gen5 NVMe	512e	SSD	7.68 TB
SK Hynix	E3.S	PCIe Gen5 NVMe	512e	SSD	7.68 TB
Kioxia	E3.S	PCIe Gen5 NVMe	512e	SSD	15.36 TB
Samsung	E3.S	PCIe Gen5 NVMe	512e	SSD	15.36 TB
SK Hynix	E3.S	PCIe Gen5 NVMe	512e	SSD	15.36 TB
Kioxia	E3.S	PCIe Gen5 NVMe	512e	SSD	1.6 TB
SK Hynix	E3.S	PCIe Gen5 NVMe	512e	SSD	1.6 TB
Kioxia	E3.S	PCIe Gen5 NVMe	512e	SSD	3.2 TB
Samsung	E3.S	PCIe Gen5 NVMe	512e	SSD	3.2 TB
SK Hynix	E3.S	PCIe Gen5 NVMe	512e	SSD	3.2 TB
Kioxia	E3.S	PCIe Gen5 NVMe	512e	SSD	3.2 TB
SK Hynix	E3.S	PCIe Gen5 NVMe	512e	SSD	3.2 TB
Kioxia	E3.S	PCIe Gen5 NVMe	512e	SSD	6.4 TB

Table 26. Supported drives (continued)

Manufacturer	Form Factor	Interface	Sector size	Drive type	Capacities
Samsung	E3.S	PCIe Gen5 NVMe	512e	SSD	6.4 TB
SK Hynix	E3.S	PCIe Gen5 NVMe	512e	SSD	6.4 TB

Solid State Drives (SSDs)

SSD Facts

Unlike hard disk drives (HDDs) which use a spinning platter to store data, solid state drives (SSDs) use solid state memory NAND flash. HDDs have several different mechanical moving parts which make them susceptible to vibrational and handling interference. Solid state drives, on the other hand have no moving parts and are less susceptible to vibrational or handling damage even when impacted during use.

SSDs deliver high-performance I/O operations per second (IOPS), and low latency for transaction - intensive server and storage applications. Properly used in systems, they reduce total cost of ownership (TCO) through low power consumption and low operating temperature.

Dell offers different solid state drive (SSD) solutions to meet different customer needs. Enterprise & Data Center SSDs, as a class, are unique compared to client or consumer-based SSD in terms of reliability, performance, and architecture. While consumer-based SSDs, such as those utilized in notebooks are designed with a focus on consumer-based workloads, rigidity and battery life, enterprise-class SSDs are designed around enterprise application I/O (I/O) requirements with focus points of random I/O performance, reliability, and protection of data during a sudden power-down.

Understanding the basics of enterprise-class SSDs allow customers to make informed decisions when comparing solutions:

- Over-provisioning: The Achilles' heel of SSDs are their write characteristics. To rewrite an area of an SSD that has already been written, the data must be erased and then written. In order to overcome a portion of the write performance penalty, Dell enterprise SSDs found across Dell PowerEdge products, all employ a practice that is known as over-provisioning of Flash. This practice keeps native Flash capacity beyond the user-defined capacity and uses the additional space as a scratch pad of sorts to quickly put down application write data on areas of Flash that are already in an erased state. The SSDs perform cleanup functions of this over-provisioned Flash space during time periods typically not impacting application performance.
- Write Endurance: Write endurance is the number of program/erase (P/E or write cycles) that can be applied to a block of flash memory before the storage media becomes unreliable. Due to different data center workloads and read/write needs, Dell offers different enterprise SSDs with different endurance ratings so customers can design the right solution for their needs.

Below are the different categories (swim lanes) of enterprise SSDs Dell offers:

- Mixed Use (MU, 3 WPD): 70/30 read/write workloads with medium endurance. E-mail/messaging, OLTP, and Ecommerce are example workloads.
- Read Intensive (RI, 1 WPD): 90/10 read/write workloads with lower endurance. Database warehousing, media streaming, and VOD solutions are example workloads.

Dell enterprise SSDs support two kinds of host interface options:

- NVMe SSD: NVMe SSDs are a mainstream, high-performance, high reliability solid-state storage device that enables IOPS performance of up to 2000x more than conventional rotating hard drives.
- SATA SSD: SATA SSDs are based on the industry-standard SATA interface. SATA SSDs provide reasonable performance for enterprise servers.

There are two classes of NVMe drives used in servers: Enterprise NVMe and Data Center NVMe SSDs:

- Data Center NVMe SSDs: This class features a balance of various factors, including performance, latency, data protection, power consumption, and affordability.
- Enterprise NVMe SSDs: Representing the premium option, this class boasts the best performance, lowest latency, robust data protection, wide capacity ranges, and extensive firmware features. However, this comes at the expense of higher power consumption and a higher price point.

Together, Dell's Enterprise and Data Center NVMe drive portfolio offers a diverse range of options for customers, covering everything from high-performance drives to cost-optimized solutions. Additionally, these drives challenge the existence of any interface other than NVMe for SSDs.

Dell Enterprise SSDs support E3.S form factor:

- E3.S: Part of the EDSFF family, E3.S is targeted to NVMe SSDs with x4 PCIe link widths. It supports power profiles up to 25 W and is positioned to be a primary form factor for mainstream NVMe server storage subsystems as it can be used across a wide variety of platforms including modular and short depth chassis.

Internal USB

An internal USB can support a small vertical USB drive with a maximum key length of 24.19 mm.

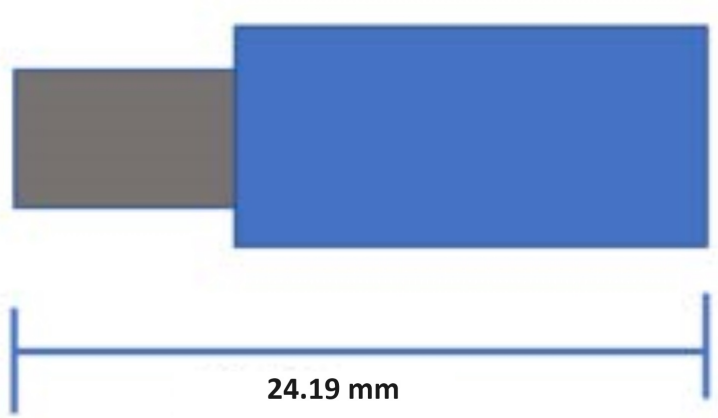

Figure 42. Internal USB

Boot Optimized Storage Solution (BOSS)

BOSS is a RAID solution that is designed to boot operating systems and separate operating system boot drives from data on server-internal storage.

BOSS feature matrix

Table 27. BOSS feature matrix

BOSS card	Drive Size	RAID levels	Stripe size	Virtual disk cache function	Maximum number of virtual disks	Maximum number of drives supported	Drive types	PCIe support	Disk cache policy	Support for Non-RAID disks	Cryptographic digital signature to verify firmware payload
BOSS-N1 Modular DC-MHS	2280 M.2 SSDs	RAID 1 and RAID 0	Supports default 64K stripe size only	None	1	2	2280 M.2 NVMe SSDs	Gen3 at 8GB/s	Drive default	No	Yes

NOTE: The system can support BOSS-N1 Modular DC-MHS on internal slot and hot-plugging of the M.2 SSDs is not supported as the as sled has to be powered off to replace the M.2 SSDs.

BOSS-N1 Modular DC-MHS

BOSS-N1 Modular DC-MHS is offered as a means of booting server to a full OS when the target OS is a full OS (not just a hypervisor), or the user does not wish to trade off standard hot plug drive slots for OS install

The BOSS-N1 Modular DC-MHS is a RAID controller with a limited feature set that presents M.2 NVMe-only SSDs as either a RAID0 disk or a single RAID1 volume with 2 disks. BOSS-N1 enables support for 480/960 GB Disks from Factory Install.

Hardware: BOSS-N1 Controller and Carrier (x2)

Reliability: Enterprise-Class M.2 NVMe SSDs

Supports dual 80 mm, Read Intensive, M.2 devices 480 GB/960 GB

Accessibility: Internal

Serviceability: Hot plug not supported

Supports Hardware RAID1 and RAID0

Supports UEFI boot

Controlled Firmware Upgrade through iDRAC

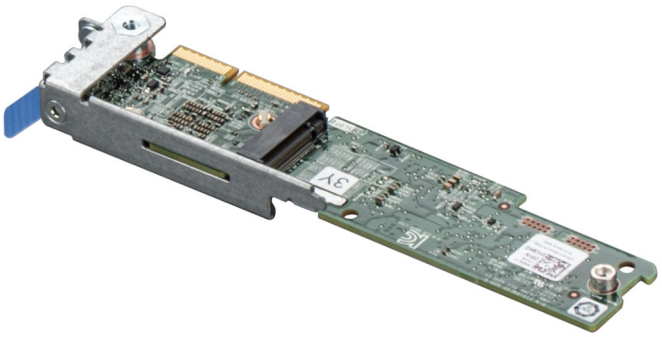

Figure 43. BOSS-N1 Modular DC-MHS

Datasheets

- BOSS-N1

BOSS User Guides

- BOSS-N1

14 Networking

Topics:

Overview

PowerEdge offers a wide variety of options to get information moving to and from our servers. Industry best technologies are chosen and these adapters are rigorously validated for worry-free, fully supported use in Dell servers.

Open Compute Project (OCP) 3.0

Supported OCP cards

OCP NIC 3.0 vs 2.0

Table 28. OCP 3.0 and 2.0 NIC comparison

Form Factor	OCP 2.0 (LOM Mezz)	OCP 3.0	Notes
PCIe Gen	Gen3	Gen5	Supported OCP3 is SFF (small form factor).
Max PCIe Lanes	Up to x16	Up to x16	See server slot priority matrix.
Shared LOM	Yes	Yes	Only OCP on slot 10 (rear IO) can support iDRAC port redirect as shared NIC.
Aux Power	Yes	Yes	Used for Shared LOM

15

PCIe subsystem

Topics:

- Expansion card installation guidelines

Expansion card installation guidelines

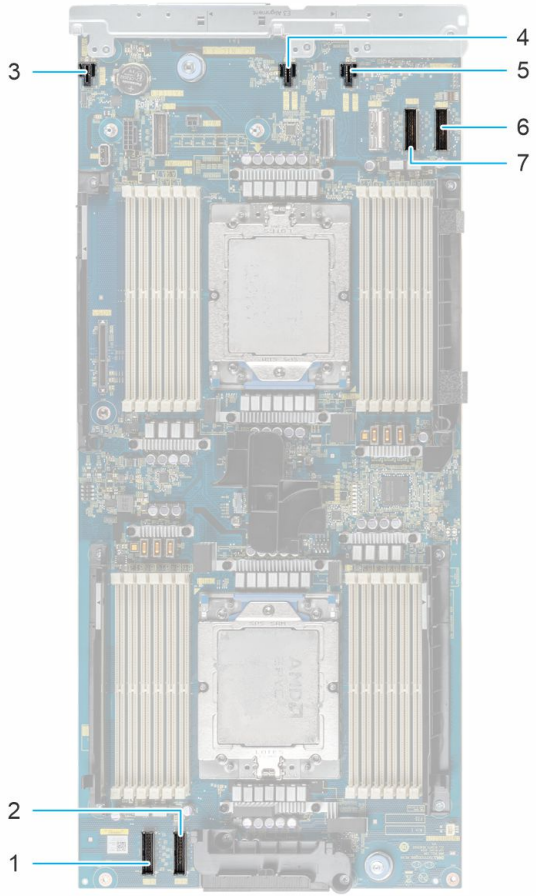

Figure 44. Expansion card riser slot connectors

1. PCIe connector 7 (SL7_CPU 1)
2. PCIe connector 6 (SL6_CPU1)
3. Riser connector (PWR 6/ PWR 7)
4. Riser connector (PWR 3/ PWR 4)
5. Riser connector (PWR 1/ PWR 2)
6. PCIe connector 1 (SL1_CPU 0)
7. PCIe connector 2 (SL2_CPU 0)

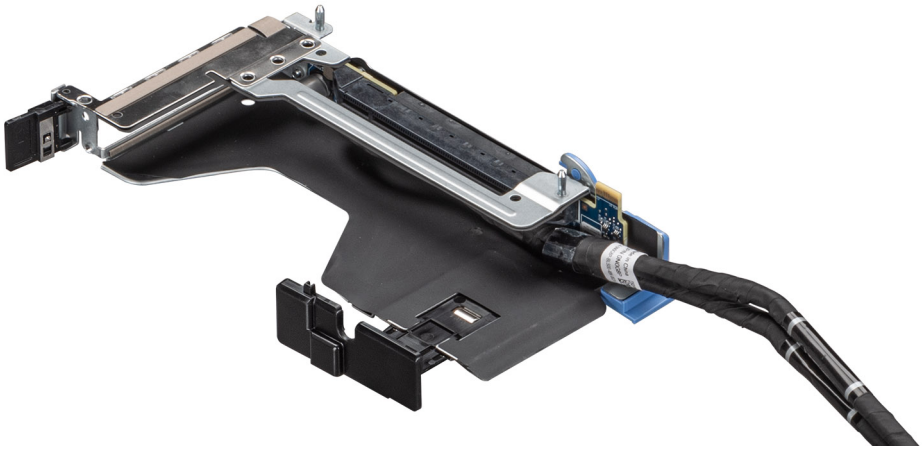

Figure 45. Riser F2A (Low Profile)

Figure 46. Riser F1A (Low Profile)

Figure 47. Riser F2B (Full Height)

Figure 48. Riser F1B (Full Height)

NOTE: The expansion-card slots are not hot-swappable.

The following table provides guidelines for installing expansion cards to ensure proper cooling and mechanical fit. The expansion cards with the highest priority should be installed first using the slot priority indicated. All the other expansion cards should be installed in the card priority and slot priority order.

The PowerEdge M7725 system supports PCI express (PCIe) slots on risers. The M7725 supports up to two Full Height, Half Lenght (FH,HL) or two Low profile PCIe slots per node.

Table 29. Expansion cards are supported on the HPM board

Config		Riser/OCP	Category	Slot	Form Factor	Slot's electrical bandwidth/ physical connector	Power
RC0	1x OCP (x16) + 2x FH (x16) per node	F1B	Cabled Riser	1	Full Height	PCIe Gen5 x16 (x16 connector)	75 W
		F2B	Cabled Riser	3	Full Height	PCIe Gen5 x16 (x16 connector)	75 W
		OCP	OCP	2	OCP	PCIe Gen5 x16 (OCP 4C+ connector)	75 W
RC1	1x OCP (x16) + 2x LP (x16) + 2x E3 (x4) per node	F1A	Cabled Riser	1	Low Profile	PCIe Gen5 x16 (x16 connector)	75 W
		F2A	Cabled Riser	3	Low Profile	PCIe Gen5 x16 (x16 connector)	75 W
		OCP	OCP	2	OCP	PCIe Gen5 x16 (OCP 4C+ connector)	75 W

16

Power, thermal, and acoustics

PowerEdge servers have an extensive collection of sensors that automatically track thermal activity, which helps to regulate temperature by reducing server noise and power consumption. The table below lists the tools and technologies Dell offers to lower power consumption and increase energy efficiency.

Topics:

Thermal

PowerEdge servers have an extensive collection of sensors that automatically track thermal activity, which helps regulate temperature thereby reducing server noise and power consumption.

Environmental specifications

Table 30. Operational Climatic Range Categories

Category A2	Allowable Operation
Temperature Ranges (For Altitude <900 meters or 2953 feet)	10 °C–35°C (50°F – 95°F) with no direct sunlight on the platform
Humidity Percent Ranges (Non-Condensing at all times)	8%RH with -12°C minimum dew point to 80%RH with 21°C (69.8°F) maximum dew point
Operational Altitude De-Rating	Maximum temperature is reduced by 1°C/300 meters (1.8°F/984 feet) above 900 meters (2,953 feet)

Table 31. Category A2

Allowable Operation	
Maximum Temperature Gradient (applies to both operation and non-operation)	20°C in an hour* (36°F in an hour) and 5°C in 15 minutes (9°F in 15 minutes), 5°C in an hour* (9°F in an hour) for tape hardware
Non-Operational Temperature Limits	-40°(-5)** C to 65°C (-40°(23) F to 149°F
Non-Operational Humidity Limits (Non-Condensing at all times)	5% to 95%RH with 27°C (80.6°F) maximum dew point.
Maximum Non-Operational Altitude	12,000 meters (39,370 feet)
Maximum Operational Altitude	3,048 meters (10,000 feet)

NOTE: *ASHARE thermal guidelines, these are not instantaneous rates of temperature change.

NOTE: **Liquid-filled components, or systems/solutions containing liquid-filled components are limited to approximately 5°C above their freeze point. At this time, the only authorized liquid coolant is Recohem PG25 with a freeze point between -9°C and -13°C, therefore the lower non-operational temperature limit is -5°C. Components and systems/solutions that can contain liquid but do not at the time of testing shall be tested to the -40°C lower non-operational temperature limit.

Table 32. Maximum vibration specifications

Maximum vibration specifications	
Operating	0.21 Grms at 5 Hz to 500 Hz (x, y, and z axes)
Storage	1.38 Grms at 7 Hz to 250 Hz for 15 min (all six sides tested)

Table 33. Maximum shock specifications

Maximum shock specifications	
Operating	Six consecutively performed shock pulses in the positive and negative x, y, and z axes of 6 G for up to 11 ms
Storage	Six consecutively performed Long Duration Half Sine Shock pulses in the positive and negative x, y, and z axes (one pulse on each side of the system) of 35 G for up to 25 ms

Particulate and gaseous contamination specifications

The following table defines the limitations that help avoid any equipment damage 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 must rectify the environmental conditions. Remediation of environmental conditions is the responsibility of the customer.

Table 34. Particulate contamination specifications

Particulate contamination	Specifications
Air filtration	Data center air filtration as defined by ISO Class 8 per ISO 14644-1 with a 95% upper confidence limit. **NOTE:** The ISO Class 8 condition applies to data center environments only. This air filtration requirement does not apply to IT equipment designed to be used outside a data center, in environments such as an office or factory floor. **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. **NOTE:** This condition applies to data center and non-data center environments.
Corrosive dust	• Air must be free of corrosive dust. • Residual dust present in the air must have a deliquescent point less than 60% relative humidity. **NOTE:** This condition applies to data center and non-data center environments.

Table 35. Gaseous contamination specifications

Gaseous contamination	Specifications
Copper coupon corrosion rate	<300A/month per class G1 as defines by ANSI/ISA71.04-2013
Silver coupon corrosion rate	<200A/month as defined by ANSI/ISA71.04-2013

Thermal design

Thermal management of the platform helps deliver high performance with the right amount of cooling to components, while maintaining the lowest fan speeds possible. This is done across a wide range of ambient temperatures from 10°C to 35°C (50°F to 95°F) and to extended ambient temperature ranges.

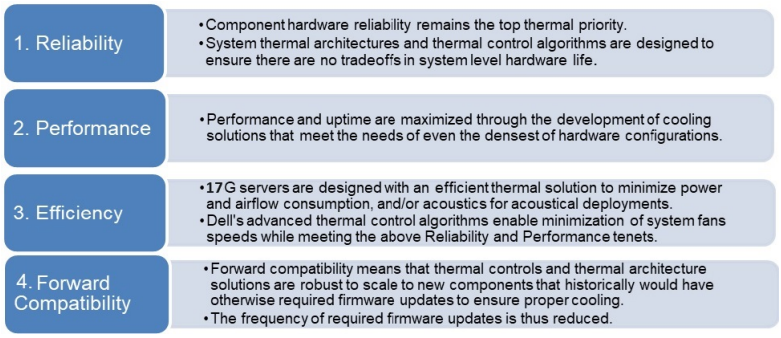

Figure 49. Thermal design characteristics

The thermal design of the PowerEdge M7725 reflects the following:

- Optimized thermal design: The system layout is architected for optimum thermal design.
- System component placement and layout are designed to provide maximum airflow coverage to critical components with minimum expense of fan power.
- Comprehensive thermal management: The thermal control system regulates the fan speed based on several different responses from all system-component temperature sensors, and inventory for system configurations. Temperature monitoring includes components such as processors, DIMMs, chipset, the inlet air ambient, hard disk drives, and OCP.
- Open and closed loop thermal fan speed control: Open loop thermal control uses system configuration to determine fan speed based on inlet air ambient temperature. Closed loop thermal control method uses feedback temperatures to dynamically determine proper fan speeds.
- User-configurable settings: With the understanding and realization that every customer has unique set of circumstances or expectations from the system. For more information, see the Dell PowerEdge M7725 Installation and Service Manual at PowerEdge Manuals and "Advanced Thermal Control: Optimizing across Environments and Power Goals" on Dell.com.
- Cooling redundancy: The M7725 allows N+1 fan redundancy, allowing continuous operation with one fan failure in the system.
- Environmental Specifications: The optimized thermal management makes the M7725 reliable under a wide range of operating environments.

Acoustics

PowerEdge acoustical specifications

Dell typically categorizes servers in five categories of acoustically acceptable usage:

- Category 1: Table-top in Office Environment
- Category 2: Floor-standing in Office Environment
- Category 3: General Use Space
- Category 4: Attended Data Center
- Category 5: Unattended Data Center

Category 1: Floor-standing in Office Environment

When Dell determines that a specific Enterprise product is to be used on a table-top in office environment, for example, on a desk around a seated user's head height, and then the acoustical specification of the following table applies. Small, light-weight towers are examples of these types of products.

Table 36. Dell Enterprise Category 1, "Table-top in Office Environment" acoustical specification category

Measurement Position re AC0158	Metric, re AC0159	Test Modes, re AC0159 (note must be in steady state, see AC0159, except where noted below)			
		Standby in 23±2° C Ambient	**Idle in 23±2° C Ambient**	**Operating in 23±2° C Ambient – if not otherwise specified in the program's configuration document, then processor and hard drive operating modes are required**	**Simulate (that is, set fan speeds representative) for Idle at 28° C & 35° C Ambient, and for 100% loading and maximum configuration, at 35° C Ambient**
Sound Power	LwA-m, bels	≤ 4.2	≤ 4.7	≤ 5.0	Report
Sound Quality (both positions must meet limits): Front Binaural HEAD and Rear Microphone	Tones, Hz, dB	No prominent tones per criteria D.10.6 and D.10.8 of ECMA-74			Report tones
	Tonality, tu	≤ 0.35	≤ 0.35	≤ 0.35	Report
	Dell Modulation, %	≤ 35	≤ 35	≤ 35	Report
	Loudness, sones	Report	Report	Report	Report
	LpA-single point, dBA	Report	Report	Report	Report
Front Binaural HEAD	Transients	• Oscillation (see AC0159), if observed, during 20-minute steady-state observation, must adhere to the following two criteria: ○ Max. {ΔLpA} < 3.0 dB ○ Event count < 3 for "1.5 dB < ΔLpA < 3.0 dB" ○ Acoustical Jump (see AC0159), during air mover speed transition from Idle to Operating Mode must be ≤ 15 dB. • Startup behavior ○ Report Startup behavior re. AC0159 ○ Startup must proceed smoothly, that is, no sudden or large jumps, and fan speed during startup must not exceed 50% of its maximum • Transient inputs: Report time-history sound pressure levels re AC0159 "Train of Step Functions on Processor"			N/A
Any	Other	No rattles, squeaks, or unexpected noises Sound should be "even" around the EUT (one side should not be dramatically louder than another) Unless otherwise specified, the "default" thermal-related settings shall be selected for BIOS and iDRAC. Specific operating conditions will be defined in "Configurations & Configuration Dependencies" for each platform.			
Sound Pressure	LpA-reported, dBA, re AC0158 and program configuration document	Report for all mics	Report for all mics	Report for all mics	Report for all mics

Category 2: Floor-standing in Office Environment

When Dell determines that a specific Enterprise product is to be used primarily when it is sitting on the floor, that is, next to a user's feet, then the acoustical specification in the table below applies. Noise from the product should not annoy or otherwise interfere with the user's thoughts or speech, for example, on the telephone.

Table 37. Dell Enterprise Category 2, "Floor-standing in Office Environment" acoustical specification category

Measurement Position re AC0158	Metric, re AC0159	Test Modes, re AC0159 (note must be in steady state, see AC0159, except where noted below)			
		Standby in 23±2° C Ambient	Idle in 23±2° C Ambient	Operating in 23±2° C Ambient – if not otherwise specified in the program's configuration document, then processor and hard drive operating modes are required	Simulate (that is, set fan speeds representative) for Idle at 28° C & 35° C Ambient, and for 100% loading and maximum configuration, at 35° C Ambient
Sound Power	LwA-m, bels	≤ 4.9	≤ 5.1	≤ 5.4	Report
Sound Quality (both positions must meet limits): Front Binaural HEAD and Rear Microphone	Tones, Hz, dB	No prominent tones per criteria D.10.6 and D.10.8 of ECMA-74			Report tones
	Tonality, tu	≤ 0.35	≤ 0.35	≤ 0.35	Report
	Dell Modulation, %	≤ 35	≤ 35	≤ 35	Report
	Loudness, sones	Report	Report	Report	Report
	LpA-single point, dBA	Report	Report	Report	Report
Front Binaural HEAD	Transients	• Oscillation (see AC0159), if observed, during 20-minute steady-state observation, must adhere to the following two criteria: ○ Max. {ΔLpA} < 3.0 dB ○ Event count < 3 for "1.5 dB < ΔLpA < 3.0 dB" • Acoustical Jump (see AC0159), during air mover speed transition from Idle to Operating Mode must be ≤ 15 dB. • Startup behavior ○ Report Startup behavior re. AC0159 ○ Startup must proceed smoothly, that is, no sudden or large jumps, and fan speed during startup must not exceed 50% of its maximum • Transient inputs: Report time-history sound pressure levels re AC0159 "Train of Step Functions on Processor"			N/A
Any	Other	• No rattles, squeaks, or unexpected noises • Sound should be "even" around the EUT (one side should not be dramatically louder than another) • Unless otherwise specified, the "default" thermal-related settings shall be selected for BIOS and iDRAC. • Specific operating conditions are defined in "Configurations and Configuration Dependencies" for each platform.			
Sound Pressure	LpA-reported, dBA, re AC0158	Report for all mics	Report for all mics	Report for all mics	Report for all mics

Table 37. Dell Enterprise Category 2, "Floor-standing in Office Environment" acoustical specification category (continued)

Measurement Position re AC0158	Metric, re AC0159	Test Modes, re AC0159 (note must be in steady state, see AC0159, except where noted below)			
		Standby in 23±2° C Ambient	Idle in 23±2° C Ambient	Operating in 23±2° C Ambient – if not otherwise specified in the program's configuration document, then processor and hard drive operating modes are required	Simulate (that is, set fan speeds representative) for Idle at 28° C & 35° C Ambient, and for 100% loading and maximum configuration, at 35° C Ambient
	and program configuration document				

Category 3: General Use Space

When Dell determines that a specific Enterprise product is to be predominantly used in a general use space, then the acoustical specification of the table below applies. These products could be found in laboratories, schools, restaurants, open office space layouts, small ventilated closets, etc., though not in close proximity to any particular person nor in quantities greater than a few in any location. People within proximity of a few of these products should not experience any impact to speech intelligibility or annoyance from the noise of the product. A rack product sitting on a table in a common area is an example.

Table 38. Dell Enterprise Category 3, "General Use" acoustical specification category

Measurement Position re AC0158	Metric, re AC0159	Test Modes, re AC0159 (note must be in steady state, see AC0159, except where noted below)			
		Standby in 23±2° C Ambient	Idle in 23±2° C Ambient	Operating in 23±2° C Ambient – if not otherwise specified in the program's configuration document, then processor and hard drive operating modes are required	Simulate (that is, set air mover speeds representative) for Idle at 28° C & 35° C Ambient and for 100% loading and maximum configuration, at 35° C Ambient
Sound Power	LwA-m, bels	≤ 5.2	≤ 5.5	≤ 5.8	Report
Sound Quality (both positions must meet limits): Front Binaural HEAD and Rear Microphone	Tones, Hz, dB	No prominent tones per criteria D.10.6 and D.10.8 of ECMA-74			Report tones
	Tonality, tu	≤ 0.35	≤ 0.35	≤ 0.35	Report
	Dell Modulation, %	≤ 40	≤ 40	≤ 40	Report
	Loudness, sones	Report	Report	Report	Report
	LpA-single point, dBA	Report	Report	Report	Report
Front Binaural HEAD	Transients	• Oscillation (see AC0159), if observed, during 20-minute steady-state observation, must adhere to the following two criteria: ○ Max. {ΔLpA} < 3.0 dB ○ Event count < 3 for "1.5 dB < ΔLpA < 3.0 dB"			N/A

Table 38. Dell Enterprise Category 3, "General Use" acoustical specification category (continued)

Measurement Position re AC0158	Metric, re AC0159	Test Modes, re AC0159 (note must be in steady state, see AC0159, except where noted below)			
		Standby in 23±2° C Ambient	**Idle in 23±2° C Ambient**	**Operating in 23±2° C Ambient – if not otherwise specified in the program's configuration document, then processor and hard drive operating modes are required**	**Simulate (that is, set air mover speeds representative) for Idle at 28° C & 35° C Ambient and for 100% loading and maximum configuration, at 35° C Ambient**

		• Report Acoustical Jump (see AC0159) during air mover speed transition from Idle to Operating Mode.
		• Startup behavior
		  ○ Report Startup behavior re. AC0159
		  ○ Startup must proceed smoothly, that is, no sudden or large jumps, and air mover speed during startup must not exceed 50% of its maximum
		• Transient inputs: Report time-history sound pressure levels re AC0159 "Train of Step Functions on Processor"			

Any	Other	No rattles, squeaks, or unexpected noises Sound should be "even" around the EUT (one side should not be dramatically louder than another) Unless otherwise specified, the "default" thermal-related settings shall be selected for BIOS and iDRAC. Specific operating conditions will be defined in "Configurations & Configuration Dependencies" for each platform.			
Sound Pressure	LpA-reported, dBA, re AC0158 and program configuration document	Report for all mics	Report for all mics	Report for all mics	Report for all mics

Category 4: Attended Data Center

When Dell determines that a specific Enterprise product is to be predominantly used in an attended data center, then the acoustical specification of the table applies. The phrase "attended data center" is used to mean a space in which many (from tens to 1000s) of Enterprise products are deployed in proximity (that is, in the same room) to personnel whose speech (perhaps with raised voices) is expected to be intelligible over the data center noise. Hearing protection or hearing monitoring programs are not expected in these areas. Examples in this category include monolithic rack products.

Table 39. Dell Enterprise Category 4, "Attended Data Center" acoustical specification category.

Measurement Position re AC0158	Metric, re AC0159	Test Modes, re AC0159 (note must be in steady state, see AC0159, except where noted below)				Simulate (that is, set fan speeds representative) for 100% loading and maximum configuration, at 35° C Ambient
		Standby in 23±2° C Ambient	Idle in 23±2° C Ambient	Operating in 23±2° C Ambient – if not otherwise specified in the program's configuration document, then processor and hard drive operating modes are required	Simulate (that is, set fan speeds representative) for Idle at 28° C & 35° C Ambient	
Sound Power	LwA-m, B	Report	≤ 6.9	≤ 7.1	Report	≤ 8.2
Front Binaural HEAD	Tones, Hz, dB	Report	< 15 dB	< 15 dB	Report	< 20 dB
	Tonality, tu	Report	Report	Report	Report	Report
	Dell Modulation, %	Report	Report	Report	Report	Report
	Loudness, sones	Report	Report	Report	Report	Report
	LpA-single point, dBA	Report	Report	Report	Report	Report
	Transients	• Oscillation (see AC0159), if observed, during 20-minute steady-state observation, must adhere to the following two criteria: ○ Max. {ΔLpA} < 3.0 dB ○ Event count < 3 for "1.5 dB < ΔLpA < 3.0 dB" ○ Acoustical Jump (see AC0159), during air mover speed transition from Idle to Operating Mode must be ≤ 15 dB. ○ Startup behavior ▪ Report Startup behavior re. AC0159 ▪ Startup must proceed smoothly, that is, no sudden or large jumps, and fan speed during startup must not exceed 50% of its maximum ∞ Transient inputs: Report time-history sound pressure levels re AC0159 "Train of Step Functions on Processor"			N/A	
Any	Other	No rattles, squeaks, or unexpected noises Sound should be "even" around the EUT (one side should not be dramatically louder than another) Unless otherwise specified, the "default" thermal-related settings shall be selected for BIOS and iDRAC. Specific operating conditions will be defined in "Configurations & Configuration Dependencies" for each platform.				
Sound Pressure	LpA-reported, dBA	Report for all mics	Report for all mics	Report for all mics	Report for all mics	Report for all mics

Category 5: Unattended Data Center

When Dell determines that a specific Enterprise product is to be predominantly used in an unattended data center (and not blades or blade enclosures; these have their own category), then the acoustical specification in the table below applies. The phrase "unattended data center" is used to mean a space in which many (from tens to 1000s) of Enterprise products are deployed together, its own heating and cooling systems condition the space, and operators or servicers of equipment enter generally only to deploy, service, or decommission equipment. Hearing protection or hearing monitoring programs may be expected (per government or company guidelines) in these areas. Examples in this category include monolithic rack products.

Table 40. Dell Enterprise Category 5, "Unattended Data Center" acoustical specification category

Measurement Position re AC0158	Metric, re AC0159	Test Modes, re AC0159 (note must be in steady state, see AC0159, except where noted below)				Simulate (that is, set air mover speeds representative) for 100% loading and maximum configuration, at 35° C Ambient
		Standby in 23±2° C Ambient	Idle in 23±2° C Ambient	Operating in 23±2° C Ambient – if not otherwise specified in the program's configuration document, then processor and hard drive operating modes are required	Simulate (that is, set air mover speeds representative) for Idle at 28° C & 35° C Ambient	
Sound Power	LwA-m, bels	Report	≤ 7.5	≤ 7.7	Report	≤ 8.7
Front Binaural HEAD	Tones, Hz, dB	Report	< 15 dB	< 15 dB	Report	< 20 dB
	Tonality, tu	Report	Report	Report	Report	Report
	Dell Modulation, %	Report	Report	Report	Report	Report
	Loudness, sones	Report	Report	Report	Report	Report
	LpA-single point, dBA	Report	Report	Report	Report	Report
Front Binaural HEAD	Transients	• Oscillation (see AC0159), if observed, during 20-minute steady-state observation, must adhere to the following two criteria: ○ Max. {ΔLpA} < 3.0 dB ○ Event count < 3 for "1.5 dB < ΔLpA < 3.0 dB" • Report Acoustical Jump (see AC0159) during air mover speed transition from Idle to Operating Mode. • Startup behavior ○ Report Startup behavior re. AC0159 ○ Startup must proceed smoothly, that is, no sudden or large jumps, and air mover speed during startup must not exceed 50% of its maximum • Transient inputs: Report time-history sound pressure levels re AC0159 "Train of Step Functions on Processor"			N/A	
Any	Other	No rattles, squeaks, or unexpected noises				

Table 40. Dell Enterprise Category 5, "Unattended Data Center" acoustical specification category (continued)

Measurement Position re AC0158	Metric, re AC0159	Test Modes, re AC0159 (note must be in steady state, see AC0159, except where noted below)				Simulate (that is, set air mover speeds representative) for 100% loading and maximum configuration, at 35° C Ambient
		Standby in 23±2° C Ambient	Idle in 23±2° C Ambient	Operating in 23±2° C Ambient – if not otherwise specified in the program's configuration document, then processor and hard drive operating modes are required	Simulate (that is, set air mover speeds representative) for Idle at 28° C & 35° C Ambient	
		Sound should be "even" around the EUT (one side should not be dramatically louder than another) Unless otherwise specified, the "default" thermal-related settings shall be selected for BIOS and iDRAC. Specific operating conditions will be defined in "Configurations & Configuration Dependencies" for each platform.				
Sound Pressure	LpA-reported, dBA, re AC0158 and program configuration document	Report for all mics	Report for all mics	Report for all mics	Report for all mics	Report for all mics

17

Rack, rails, and cable management

Topics:

Rack, rails and cable management

The rail offerings for the PowerEdge M7725 consist of static rails. The cable management offerings, cable management arm (CMA) and an optional strain relief bar (SRB) are not supported.

Static rails Information

The M7725 support L-bracket static rails. The L-bracket static rails (shown in the figure below) matches Open Rack Frame V3 specification, do not support serviceability in the rack.

L-Bracket Static Rails for 21-inch Open Computing Project Racks:

- Snap into the sides of the rack.
- Support installation in 2-inch ORv3 rack

Figure 50. Static rails for M7725

Rack installation

The M7725 offers L-bracket rail that matched Open Rack Frame V3 specification and snap into the sides of ORv3 Rack.

Figure 51. Installing the static rails in IR7000 rack before installing M7725 system

18

Operating Systems and Virtualization

Topics:

- Supported operating systems

Supported operating systems

The Dell PowerEdge M7725 supports the following operating systems:

- Canonical Ubuntu Server LTS
- SUSE Linux Enterprise Server
- Red Hat Enterprise Linux

For specifications and interoperability details, see OS support. .

19

Dell Systems Management

Dell delivers management solutions that help IT administrators deploy, update, monitor, and manage IT assets. OpenManage solutions and tools enable you to solve and respond to problems quickly by manage Dell servers efficiently in physical and remote environments, and operating in-band and out-of-band (agent-free).

The OpenManage portfolio includes innovative embedded management tools such as the integrated Dell Remote Access Controller (iDRAC) and consoles like OpenManage Enterprise, OpenManage Power Manager Plugin, and tools like Repository Manager. Dell has developed comprehensive systems management solutions that are based on open standards by connecting and/or integrating it's offers with top system management vendors and frameworks such as Ansible, Microsoft, and VMware, enabling advanced management of Dell hardware. The key tools for managing Dell PowerEdge servers are iDRAC and OpenManage Enterprise (OME) console. OpenManage Enterprise helps the system administrators with the life cycle management of multiple generations of PowerEdge servers. OME has additional functions that can be added with plugins like OpenManage Enterprise Services, Update Manager, APEX AIOps Observability (formerly CloudIQ), and Power Manager. It also offers integration with VMware vCenter and Microsoft System Center, and a set of tools, including Repository Manager, enabling easy management of PowerEdge hardware. The four main pillars of Dell systems management closely align with the issues and business challenges that are faced by many IT departments.

- Automating IT management.
  - Comprehensive automation management for reducing OPEX and increasing uptime and overall efficiency of systems.
  - Comprehensive suite of tools to automate according to your needs.
- Management made simple.
  - Simple but powerful tools for managing your Dell servers.
  - Integrated tools that streamline support engagements.
  - Innovative at-the-box management features.
- Secure by default.
  - Dell servers offer robust security defenses to prevent the next generation of malicious attacks.
  - Security is designed deep into the hardware and firmware architecture for optimal protection.
- Smarter infrastructure management.
  - It offers a next-generation 1-to-many console to manage your IT and server infrastructure.
  - Embedded intelligence that is infrastructure-aware to optimize troubleshooting and deployment.

This document provides an overview of the OpenManage Systems Management offerings to help IT administrators choose the appropriate tools to completely manage Dell PowerEdge servers.

- The latest Dell Systems Management Overview Guide.

Topics:

Integrated Dell Remote Access Controller (iDRAC)

iDRAC10 delivers advanced, agent-free, local and remote server administration. Embedded in every PowerEdge server, iDRAC10 provides a secure means to automate a multitude of common management tasks. Because iDRAC is embedded within every PowerEdge server, there is no additional software to install; plug in power and network cables, and iDRAC is ready to go. Even before installing an operating system (operating system) or hypervisor, IT administrators have a complete set of server management features at their fingertips.

With iDRAC10 in-place across the Dell PowerEdge portfolio, the same IT administration techniques and tools can be applied throughout. This consistent management platform allows scaling of PowerEdge servers as an organization's infrastructure grows. Customers can use the iDRAC RESTful API for the latest in scalable administration methods of PowerEdge servers. With this API, iDRAC enables support for the Redfish standard and enhances it with Dell extensions to optimize at-scale management of PowerEdge servers.

Zero-Touch Provisioning (ZTP) is embedded in iDRAC. ZTP is an Intelligent Automation Dell's agent-free management. Once a PowerEdge server is connected to power and networking that system can be monitored and fully managed, whether you are standing in front of the server or remotely over a network. With no need for software agents, an IT administrator can:

- Monitor
- Manage
- Update
- Troubleshoot, and remediate Dell servers.

With features like zero-touch deployment and provisioning, and System Lockdown, iDRAC10 is purpose-built to simplify server administration. For those customers whose existing management platform uses in-band management, Dell does provide iDRAC Service Module, a lightweight service that can interact with both iDRAC10 and the host operating system to support legacy management platforms.

When ordered with DHCP enabled from the factory, PowerEdge servers can be automatically configured when they are initially powered up and connected to your network. This process uses profile-based configurations that ensure each server is configured per your specifications. This feature requires an iDRAC Enterprise license.

iDRAC10 offers the following license tiers:

Table 41. iDRAC10 license tiers

License	Description
iDRAC10 Core	• Available for all servers. • Core system management features for users who are cost conscious.
iDRAC10 Enterprise	• Available as an upsell on all servers. • Includes all features of Core. Also, includes additional automation features and virtual console and security features. • Bundled with Secure Enterprise Key Management (SEKM) and Secure Component Verification (SCV) licenses.
iDRAC10 Datacenter	• Available as an upsell on all servers. • Includes all features of Core and Enterprise. • Includes key features such as telemetry streaming and thermal management. • Includes advanced accelerators (GPU and DPU) system management and advanced air and liquid cooling.

For a full list of iDRAC features by license tier, see the **Integrated Dell Remote Access Controller 10 User's Guide** at Dell.com.

For more details on iDRAC10 including white papers and videos, see:

- Support for Integrated Dell Remote Access Controller 10 (iDRAC10) is on the Knowledge Base page at Dell.com

Systems Management software support matrix

Table 42. Systems Management software support matrix

Categories	Features	PE mainstream
Embedded Management	iDRAC10	Supported
	iDRAC Direct	Supported
	iDRAC RESTful API with Redfish	Supported
Change Management	Dell Repository Manager	Supported
	Dell System Update	Supported
	Enterprise Catalogs	Supported
	Server Update Utility (SUU)	Supported
OpenManage console	CloudIQ for PowerEdge plug-in	Supported
	OpenManage Enterprise (OME)	Supported

Table 42. Systems Management software support matrix (continued)

Categories	Features	PE mainstream
	OME APEX AIOps Observability	Supported
	OME integration for Microsoft System Center	Supported
	OME Integration for VMware vCenter (with VMware Aria Operations)	Supported
	OpenManage Integration for Windows Admin Center	Supported
	OME Power Manager	Supported
	OME Services	Supported
	OME Update Manager	Supported
OpenManage Integrations	BMC TrueSight	Supported
	Microsoft System Center	Supported
	OpenManage Integration with ServiceNow	Supported
	Red Hat Ansible Modules	Supported
	Terraform Providers	Supported
	VMware vCenter and vRealize Operations Manager	Supported
Security	Cryptographically signed firmware	Supported
	Secure Boot	Supported
	Secured Component Verification (Hardware integrity check)	Supported
	Secure Erase	Supported
	Silicon Root of Trust	Supported
	TPM 2.0 FIPS, CC-TCG certified	Supported
	AMD Secure Memory Encryption (SME)	Supported
	AMD Secure Encrypted Virtualization (SEV)	Supported
Operating system	Canonical Ubuntu Server LTS	Supported

20

Appendix D: Service and support

Topics:

Why attach service contracts

Dell PowerEdge servers include a standard hardware warranty that highlights our commitment to product quality by guaranteeing repair or replacement of defective components. While industry-leading, our warranties are limited to 1 or 3 years, depending on model, and do not cover software assistance. Call records show that customers are most often seeking Dell technical support for software related issues like configuration guidance, troubleshooting, upgrade assistance or performance tuning. Encourage customers to purchase ProSupport service contracts to supplement warranty coverage and ensure optimal support for both hardware and software. ProSupport provides a complete hardware guarantee beyond the original warranty period.

ProSupport Infrastructure Suite

ProSupport Infrastructure Suite is a set of support services that enable customers to build the solution that is right for their organization. It is an industry-leading, enterprise-class support that aligns with the criticality of your systems, the complexity of your environment, and the allocation of your IT resources.

Figure 52. ProSupport Enterprise Suite

	Basic Hardware Support[1]	ProSupport	BEST ProSupport Plus
Outcome Assistance and Advocacy *via assigned Technical Customer Success Manager* ⓘ			
Enjoy a frictionless customer experience with cross-functional lifecycle management aligned to your goals			✓
Accelerate time-to-value through onboarding assistance, education and success planning			✓
Turn challenges into opportunities with actionable strategies powered by data and AI-driven analytics			✓
Ensure coverage continuity while preparing to scale for future success			✓
Proactive Monitoring & Actionable Insights *via Dell's connectivity solutions and tools*			
Quickly visualize performance through a current system health score		✓	✓
Cybersecurity monitoring and mitigation recommendations provide another layer of protection		✓	✓
Predictive performance and capacity analysis address bottlenecks		✓	✓
Prevent or plan for downtime with predictive hardware anomaly detection		✓	✓
Energy consumption and carbon footprint forecasting support sustainability and stewardship initiatives		✓	✓
Get ahead of problems with proactive issue detection with automated case creation	✓	✓	✓
Streamline internal IT efforts with efficient service request and escalation management tools	✓	✓	✓
Minimize disruptions by self-dispatching eligible parts	✓	✓	✓
Support Essentials			
Receive an assigned incident manager for Sev 1 issues who will work your issue through to resolution		✓	✓
Count on Mission Critical Support during Sev 1 incidents and natural disasters ⓘ			✓
Keep systems code current and performing at peak through Proactive System Maintenance			✓
Get priority access to senior technical support engineers—skip the queues and callbacks			✓
Bringing your own software? We provide limited 3rd party software support ⓘ			✓
Choose onsite parts delivery and labor response that meets your needs	Next Business Day	NBD or 4-hour	4-hour
Select product coverage that best augments your internal resources	Hardware	Hardware & Software	Hardware & Software
Have an issue? We are here for you by phone, chat and online	Local business hours	24/7/365	24/7/365

ProSupport Plus for Infrastructure

ProSupport Plus for Infrastructure is the ultimate solution for customers seeking preventative maintenance and optimal performance on their business-critical assets. The service caters to customers who require proactive, predictive, and personalized support for systems that manage critical business applications and workloads. When customers purchase PowerEdge server, we recommend ProSupport Plus, our proactive and preventative support service for business-critical systems. ProSupport Plus provides all the benefits of ProSupport, including the following "Top five reasons to buy ProSupport Plus (PSP)"

1. **Priority access to specialized support experts:** Immediate advanced troubleshooting from an engineer that understands Dell infrastructure solutions.
2. **Mission Critical Support:** When critical (Severity 1) support issues happen, the customer is assured that we do all that we can to get them back up and running as quickly as possible.
3. **Technical Customer Success Manager:** A customer's #1 support advocate, ensuring they get the best possible proactive and predictive support experience.
4. **Systems maintenance:** On a semiannual basis, we will keep a customer's ProSupport Plus system(s) up to date by installing the latest firmware, BIOS, and driver updates to improve performance and availability.
5. **Third-party software support:** Dell is a customer's single point of accountability for any eligible third-party software that is installed on their ProSupport Plus system, whether they purchased the software from us or not.

ProSupport for Infrastructure

Comprehensive 24x7 support for hardware and software – best for production, but not critical, workloads and applications. The ProSupport service offers highly trained experts around the clock and around the globe to address IT needs. We help minimize disruptions and maximize availability of PowerEdge server workloads with:

- 24x7 support through phone, chat and online
- A central point of accountability for all hardware and software issues
- Hypervisor, operating system, and application support
- Dell security advisories

- Onsite response service levels 4 hour or Next Business Day options
- Proactive issue detection with automated case creation
- Predictive hardware anomaly detection
- Incident Manager assigned for Severity 1 cases
- Collaborative third-party support
- Access to AIOps Platforms - (MyService360, TechDirect, and CloudIQ)
- Consistent experience regardless of where customers are located or what language that they speak.

Basic Hardware Support

Provides reactive hardware support during normal business hours, excluding local national holidays. No software support or software-related guidance. For improved levels of support, choose ProSupport or ProSupport Plus.

Specialty Support Services

Optional specialty support services complement the ProSupport Infrastructure Suite to provide additional proficiencies that are critical for modern data center operations.

Hardware coverage add-ons to ProSupport or ProSupport Plus

- **Keep Your Hard Drive (KYHD), Keep Your Component (KYC), or Keep Your GPU (KYGPU):**

  Normally if a device fails under warranty, Dell replaces it using a one-for-one exchange process.KYHD/KYCC/KYGPU gives you the option to retain your device. It provides full control of sensitive data and minimizes security risk by letting you retain possession of failed drives, components, or GPU when receiving replacement parts without incurring additional cost.
- **Onsite Diagnosis Service:**

  Ideal for sites with non-technical staff. Dell field technician performs initial troubleshooting diagnosis onsite and transfers to Dell remote engineers to resolve the issue.
- **ProSupport Add-on for HPC:**

  Sold as an add-on to a ProSupport service contract, the ProSupport Add-on for HPC provides solution-aware support to cover the additional requirements that are required to maintain an HPC environment such as:

  - Access to senior HPC experts
  - Advanced HPC cluster assistance: Performance, interoperability, and configuration
  - Enhanced HPC solution level end-to-end support
  - Remote pre-support engagement with HPC Specialists during ProDeploy implementation
- **ProSupport Add-on for Telco (Respond & Restore):**

  An add-on service designed for the top 31 TELCO customers globally, Respond & Restore provides direct access to Dell solution experts who specialize in TELCO carrier-grade support. This add-on also provides a hardware uptime guarantee, meaning if a system fails, Dell has it installed and operational within 4 hours for Severity 1 issues. Dell incurs penalties and fees if SLAs are not met.

Personalized Support and Supplemental Site-wide Expertise

- **Technical Account Manager:**

  Designated technology lead who monitors and manages the performance and configuration of specific technology sets.
- **Designated Remote Support:**

  Personalized support expert who manages all troubleshooting and resolution of IT assets.
- **Multivendor Support Service:**

  Support your third-party devices as one service plan for servers, storage, and networking (includes coverage for: Broadcom, Cisco, Fujitsu, HPE, Hitachi, Huawei, IBM, Lenovo, NetApp, Oracle, Quanta, SuperMicro and others).

Services for large enterprises

- **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 (combined total of server, storage, networking, so forth). This offering is built on standard ProSupport features that leverage our global scale and are tailored to specific customer needs. While not for everyone, this service option offers a truly unique solution for our largest customers with the most complex environments.

  - Team of assigned Services Account Managers with remote or onsite options
  - Assigned technical and field engineers who are trained on the customer's environment and configurations.
  - On-demand reporting and recommendations that are enabled by ProSupport AIOps tools (MyService360, TechDirect, and CloudIQ)
  - Flexible onsite support and parts options that fit their operational model
  - A tailored support plan and training for their operations staff
- **ProSupport One for Data Center – CSP (Cloud Serviced Provider) and AI Solution**

  ProSupport One for Data Center – CSP and AI Solution is a unique offer that is designed for a limited set of Dell accounts purchasing AI computing solutions greater than 1,000 servers and $250M in sales. PS1DC - CSP and AI improves the entire services experience combining support, deployment (rack integration), residency services, a designated support engineer, an onsite service engineer, and an onsite parts service as one holistic offer. Special pricing has been determined to compete effectively against competitors and provide the best customer experience. PS1DC for CSP and AI can only be sold with XE Servers and all networking platforms (Dell and NVIDIA). All other products would be eligible for the standard PS1DC, not this unique offer. More details on PS1DC for CSP and AI here.
- **Onsite Parts Service (OPS)**

  Ideal for large organizations that have their own staff to support their data center. Dell offers a service that is called Onsite Parts Service (OPS) from Dell Services. OPS manages parts inventory located at the customer's designated facility. The Logistics Online Inventory Solution (LOIS) program will use software to support the monitoring and automatic replenishment of inventory stored on the customer site. . Each replacement part would automatically initiate a replenishment of the parts inventory that is shipped the next day or delivered onsite by Dell during a regular scheduled visit (called Scheduled Onsite Service). As part of the LOIS system, customers can integrate their systems directly to Dell TechDirect using APIs to help streamline the support management process.

End-of-Life Services

- **Post Standard Support (PSS)**

  Extend service life beyond the initial seven years of ProSupport, adding up to five more additional years of hardware coverage.
- **Data Sanitization & Data Destruction**

  Renders data unrecoverable on repurposed or retired products, ensuring security of sensitive data and enabling compliance and provides NIST-compliant certification.
- **Asset Recovery Services**

  Recycle, resale, and disposal of hardware. Helps you securely and responsibly retire IT assets that are no longer needed while protecting both your business and the planet.

ProDeploy Infrastructure Suite

ProDeploy Infrastructure Suite provides various deployment offerings that satisfy a customer's unique needs. It is made up of various sub-offers: Factory Configuration Services, Rack Integration, Basic Deployment, ProDeploy, ProDeploy Plus, and optionally ProDeploy FLEX which allows for some customization of the features listed.

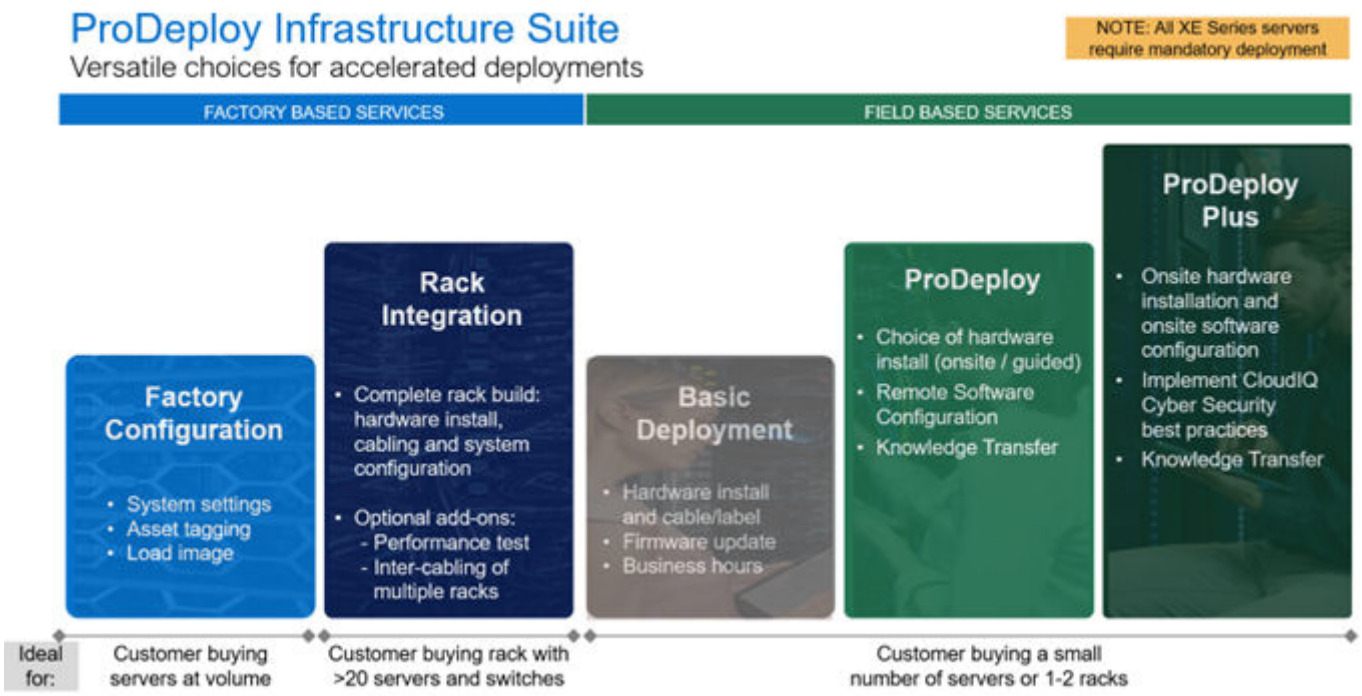

Figure 53. ProDeploy Infrastructure Suite

Factory-based Services

Pre-configured systems or complete racks, customized prior to shipping to the customer's site.

Customer Rack Integration or ProDeploy FLEX Rack Integration

Dell offers robust custom rack integration services through two main programs: Enterprise Rack Integration Services and Integrated Rack Scalable Systems (IRSS). These services are designed to streamline deployment, reduce complexity, and optimize performance for data centers, edge environments, and AI workloads. These factory services are purchased as a custom engagement or as ProDeploy Flex Rack Integration SKUs.

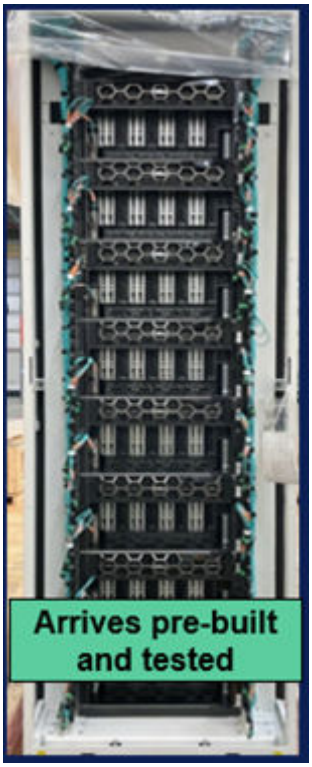

Figure 54. Pre-configured system

Figure 55. Pre-configured system

Factory Configuration

Ideal for customers buying servers in volume and seeking pre-configuration prior to shipping such as: custom image, system settings, and asset tagging so it arrives ready to use out of the box. Furthermore, servers are packaged and bundled to meet specific shipping and distribution requirements for each customer location to facilitate the rollout process. Once the server is onsite, Dell can install and configure the server to the environment using any of the field-based deployment services outlined in the next section.

Field-based services

Put PowerEdge servers to work faster with Dell field-based deployment services. Whether we are deploying one server to one thousand – we have you covered. Dell provides versatile delivery options to fit every budget and operating model.

- **ProDeploy Plus:** Elevate Infrastructure deployments with our most complete service from planning through onsite hardware installation and software configuration including the implementation of cybersecurity best practices. ProDeploy Plus provides the skill and scale that is needed to successfully execute demanding deployments in today's complex IT environments. The deployment starts with a site readiness review and implementation plan. Certified deployment experts perform the software configuration to include setup of leading operating systems and hypervisors. Dell will also configure PowerEdge software tools to include iDRAC and OpenManage system utilities as well as support AIOps platforms: MyService360, TechDirect, and

CloudIQ. Unique to ProDeploy Plus, the cybersecurity implementation helps customers understand potential security risks and make recommendations for reducing product attack surfaces. The system is tested, validated prior to completion. The customer will also receive full project documentation and knowledge transfer to complete the process.

- **ProDeploy:** ProDeploy provides remote software configuration and choice of hardware installation (onsite or guided). ProDeploy is great for customers who are price sensitive or willing to participate in some portion of the deployment to include providing remote access to their network. The ProDeploy remote software includes everything mentioned in ProDeploy Plus except it does not include the added value, cybersecurity implementation, and implementatiod best practices.
- **Basic Deployment:** Basic Deployment delivers worry-free professional installation by experienced technicians. This service is often sold to Competency Enabled Partners who will have Dell do the hardware installation while they complete the software configuration. Furthermore, Basic Deployment tends to be purchased by large enterprises who have smart technical staff. These companies just need Dell to install the hardware, and they will perform the software configuration. The last use case for Basic Deployment is when paired with Factory Configuration services. The servers are preconfigured in the factory, and the basic deployment service will install the system into the rack to finalize the deployment.

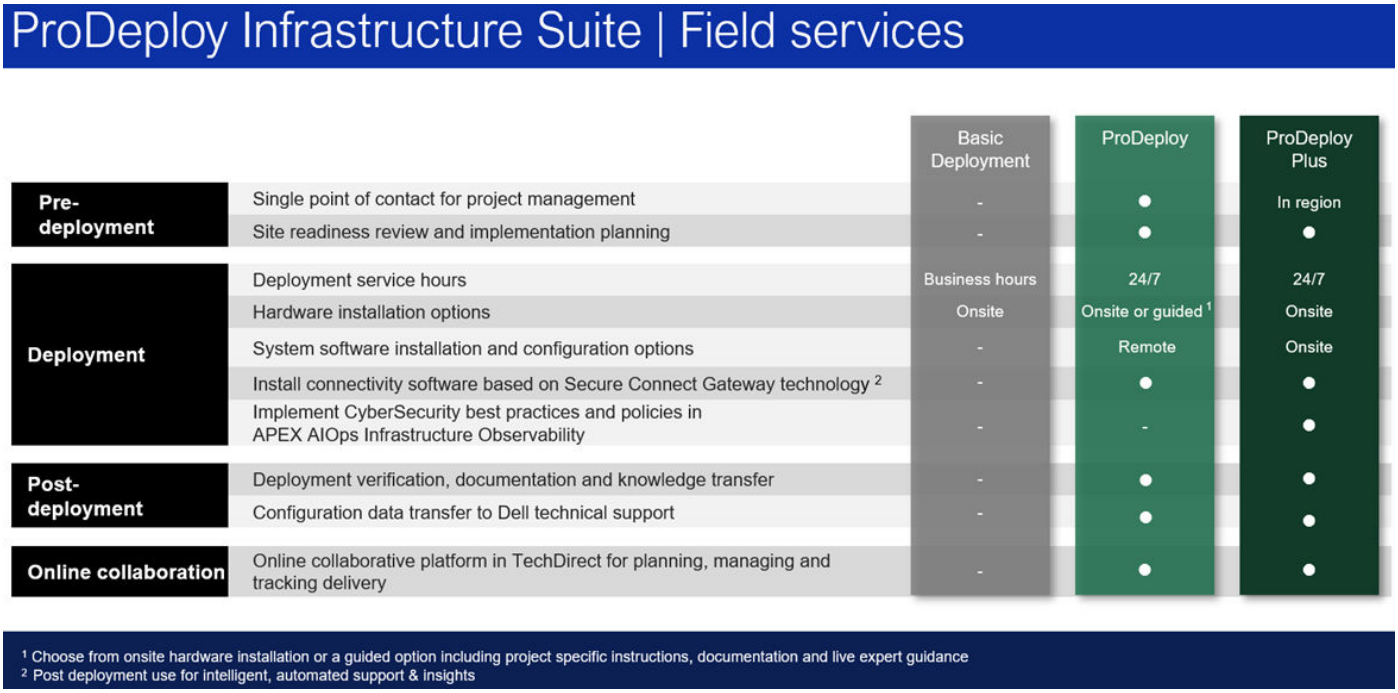

Figure 56. ProDeploy Infrastructure Suite - Field services

Supplemental Deployment Services

Additional ways to expand scope or deploy for unique scenarios.

Two Host Adder (requires PD/PDP)

Deploying new storage, compute, or networking devices may require interconnection to other servers (also called hosts). The Dell delivery team will set up four hosts per device as part of every ProDeploy service. For example, if the customer is buying two storage arrays the ProDeploy service will automatically include connectivity of four hosts each (4x2=8 total hosts per project since there are two devices). This supplemental “Two Host Adder” service provides for the configuration of additional hosts above what is already provided as part of the ProDeploy service. In many cases, customers can work with us while we set up the included hosts, so they may understand how to do the rest themselves. Always ask the customer how many hosts are being connected and sell the host adder depending on the customer’s technology skillset. Note that this service applies to the connectivity of Dell devices not 3rd party devices.

Additional Deployment Services (ADT) – sold with or without PD/PDP

You can expand the scope of a ProDeploy engagement leveraging Additional Deployment Time (ADT). ADT covers additional tasks above the normal deliverables of the ProDeploy offers. ADT can also be used as a standalone service without ProDeploy.

SKUs are available for both Project Management and Technical Resource Expertise. SKUs are sold as blocks of four hours remote or eight hours onsite. The delivery team can help in scoping the number of hours required for additional tasks.

Data Migration Services

Migrating data sets is no easy task. Our experts use proven tools and process to streamline data migrations and avoid compromising data. A customer project manager works with our experienced team of experts to create a migration plan. Data migration is part of every technology upgrade, platform change, and shift to the cloud. You can rely on Dell data migration services to perform a seamless transition.

Residency Services

Certified technical professionals act like an extension of your IT staff to enhance internal capabilities and resources and help you realize faster adoption and maximized ROI of new technology. Residency Services help customers transition to new capabilities quickly by leveraging specific technology skill sets. Residency experts can provide post implementation management and knowledge transfer that is related to a new technology acquisition or day-to-day operational management of the IT infrastructure.

- Global experts available to serve in-person (onsite) or virtual (remote)
- Engagements starting at 2 weeks with flexibility to adjust
- Residency is available for project management needs, and many different technology skills sets such as: Server, storage, Gen AI, networking, security, multi-cloud, data mgmt., and modern workforce application residents

Unique Deployment Scenarios

Custom Deployment Services

When a deployment is beyond the scope of the ProDeploy Infrastructure Suite, you can turn to the custom deployment services team to address complex implementation scenarios and customer unique requirements. The Dell custom deployment team is staffed with solution architects who will assist with customer scoping calls to define the project and develop the statement of work. Custom services can handle a wide range of deployments that can be performed in the factory or onsite. All custom engagement services are requested through SFDC.

Deployment of AI or HPC using Cluster Build Services

Once the integrated rack arrives the data center or is built onsite, Dell can also convert the racks into a large computing cluster. Dell provides several deploy options for Artificial Intelligence (AI) or High-Performance Computing (HPC) implementations. These complex environments require specialists that understand advanced feature sets to create a unified computing cluster for the most demanding workloads. Choose one of the cluster build add-ons below.

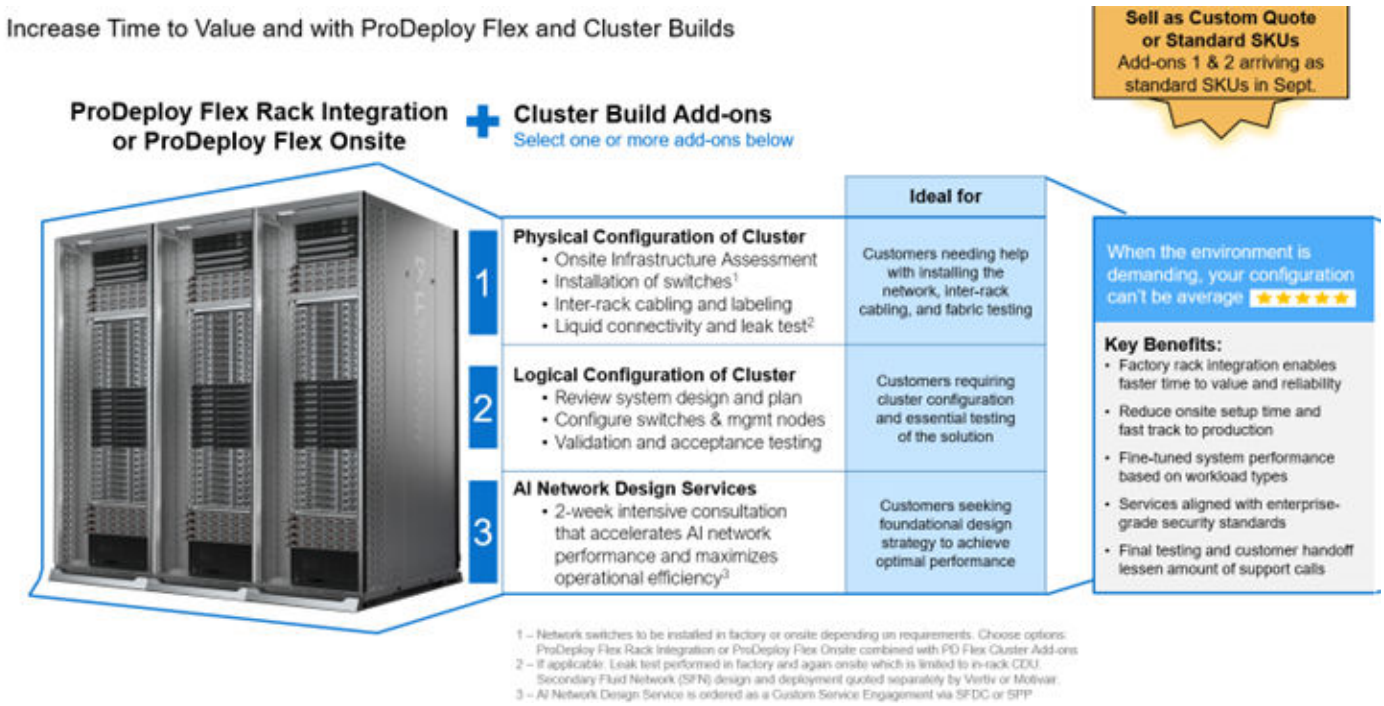

Figure 57. Deployment choices for cluster implementation

DAY 2 – Automation Services with Ansible

Dell solutions are built as "automation ready" with integrated APIs (Application Programming Interfaces) to allow customers to programmatically call actions on the product through code. Although Dell has published Anisble automation use cases, some customers need additional assistance with GitOps. By the end of the service, the customer will have the foundational components required to accelerate automation and understand how the programming works together: Day 1 and Day 2 use case automation scripts (ansible modules), CI/CD tool (Jenkins), and Version control (Git).

Dell Technologies Consulting Services

Our expert consultants help customers transform faster, and quickly achieve business outcomes for the high-value workloads Dell PowerEdge systems can handle. From strategy to full-scale implementation, Dell Technologies Consulting can help determine how to perform IT, workforce, or application transformation. We use prescriptive approaches and proven methodologies that are combined with the portfolio and partner ecosystem of Dell Technologies to help achieve real business outcomes. From multicloud, applications, DevOps, and infrastructure transformations, to business resiliency, data center modernization, analytics, workforce collaboration, and user experiences-we are here to help.

Dell Managed Services

Some customers prefer Dell to manage the complexity and risk of daily IT operations, Dell Managed Services utilizes proactive, AI enabled delivery operations and modern automation to help customers realize desired business outcomes from their infrastructure investments. With these technologies, our experts run, update, and fine-tune customer environments that are aligned with service levels, while providing environment-wide and down-to-the-device visibility. There are two types of managed service offers. First the outsourcing model or CAPEX model where Dell manages the customer owned assets using our people and tools. The second is the as-a-Service model or OPEX model called APEX. In this service, Dell owns all technology and all the management of it. Many customers will have a blend of the two management types depending on the goals of the organization.

Figure 58. Dell Managed Services

Cyber-Security Services

Managed Detection and Response (MDR)

Dell Managed Detection and Response Pro Plus is our fully-managed, 360º security operations solution comprised of our most cutting-edge, preventive and responsive cybersecurity services. MDR Pro Plus was designed with your top security concerns in mind, allowing you to focus on your core business goals while Dell handles your security operations. First, we have Vulnerability Management. With this service, we'll do ongoing scanning of the customer's environment looking for software that needs to be patched. Next is Pen Testing and Attack Simulation Management. This service will continuously validate security controls and policies with automated Breach and Attack Simulation (BAS), because a misconfiguration can lead to an exposure which an attacker can exploit. The service also includes an annual penetration test to determine if a skilled threat actor could exploit pathways leading to critical assets or data. Third, Managed Security Awareness Training. This service will educate the customer's end users so that they don't inadvertently put the customer at risk. If you think about our annual compliance training modules, there is always a security module. This is the same type of thing, but rather than once a year, it will be smaller, bite-size pieces of content delivered throughout the year.Fourth is our Managed Detection and Response service which provides 24x7 threat detection and investigation, analysis of end-to-end activity by threat actors, threat hunting, and quick initiation of cyber incident response when needed. Customers can choose between Secureworks Taegis XDR, CrowdStrike Falcon XDR or Microsoft Defender XDR as the security analytics platform our analysts will use to monitor their environment. All four of these services are delivered by experienced, certified Dell security experts using advanced technology such as the Secureworks Taegis XDR, CrowdStrike Falcon XDR or Microsoft Defender XDR security platforms.

Dell Technologies Education Services

Build the IT skills required to influence the transformational outcomes of the business. Enable talent and empower teams with the right skills to lead and perform transformational strategy that drives competitive advantage. Leverage the training and certification that is required for real transformation.

Dell Technologies Education Services offers PowerEdge server training and certifications that are designed to help customers achieve more from their hardware investment. The curriculum delivers the information and the practical, firsthand skills that their team must confidently install, configure, manage, and troubleshoot Dell servers.

To learn more or register for a class today, see Education.Dell.com.

Resources

Service for powerEdge

21

Appendix B. Standards compliance

The system conforms to the following industry standards.

Table 43. Industry standard documents

Standard	URL for information and specifications
ACPIAdvance Configuration and Power Interface Specification, v6.4	ACPI
Ethernet IEEE Std 802.3-2022	IEEE Standards
MSFT WHQL Microsoft Windows Hardware Quality Labs	Windows Hardware Compatibility Program
IPMI Intelligent Platform Management Interface, v2.0	IPMI
DDR5 Memory DDR5 SDRAM Specification	DDR5 SDRAM
PCI Express PCI Express Base Specification, v5.0	PCIe specifications
PMBus Power System Management Protocol Specification, v1.2	PMBus specifications
SMBIOS System Management BIOS Reference Specification, v3.3.0	DMTF SMBIOS
TPM Trusted Platform Module Specification, v2.0	TPM specifications
UEFI Unified Extensible Firmware Interface Specification, v2.7	UEFI specifications
PI Platform Initialization Specification, v1.7	
USB Universal Serial Bus v2.0 and SuperSpeed v3.0 (USB 3.1 Gen1)	USB document library
NVMe Express Base Specification. Revision 2.0c	NVMe specifications
NVMe Command Set Specifications 1. NVM Express NVM Command Set Specification. Revision 1.1c 2. NVM Express Zoned Namespaces Command Set. Revision 1.0c 3. NVM Express® Key Value Command Set. Revision 1.0c	
NVMe Transport Specifications 1. NVM Express over PCIe Transport. Revision 1.0c 2. NVM Express RDMA Transport Revision. 1.0b 3. NVM Express TCP Transport. Revision 1.0c	
NVMe NVM Express Management Interface. Revision 1.2c	
NVMe NVMe Boot Specification. Revision 1.0	

22

Appendix C: Additional resources

Table 44. Additional resources

Resource	Description of contents	Location
Installation and Service Manual	This manual, available in PDF format, provides the following information: • Chassis features • System Setup program • System indicator codes • System BIOS • Remove and replace procedures • Diagnostics • Jumpers and connectors	Dell.com/Support/Manuals
Getting Started Guide	This guide ships with the system, and is also available in PDF format. This guide provides the following information: • Initial setup steps	Dell.com/Support/Manuals
Rack Installation Guide	This document ships with the rack kits, and provides instructions for installing a server in a rack.	Dell.com/Support/Manuals
System Information Label	The system information label documents the HPM board layout and system jumper settings. Text is minimized due to space limitations and translation considerations. The label size is standardized across platforms.	Inside the system chassis cover
MyDell label	This code on the chassis can be scanned by a phone application to access additional information and resources for the server, including videos, reference materials, service tag information, and Dell contact information.	Inside the system chassis cover
Enterprise Infrastructure Planning Tool (EIPT)	The Dell online EIPT enables easier and more meaningful estimates to help you determine the most efficient configuration possible. Use EIPT to calculate the power consumption of your hardware, power infrastructure, and storage.	Dell.com/calc