2

System features and generational comparison

The following table shows the comparison between the PowerEdge R960 with the PowerEdge R940xa.

Table 2. Features comparison

Features	PowerEdge R960	PowerEdge R940xa
Processors	4 x 4th Gen Intel® Xeon® Scalable Processors	4 x 2nd Generation Intel® Xeon® Scalable Processor
CPU interconnect	Intel Ultra Path Interconnect (UPI)	Intel Ultra Path Interconnect (UPI)
Memory	• 64 x DDR5 RDIMM • Up to 4800 MT/s (1 DPC) / 4400 MT/s (2 DPC)	• 48 x DDR4 RDIMM, LRDIMM • 24 x PMem (Intel Optane Persistent Memory) • 12 x NVDIMM
Storage Controllers	• PERC 11G: H755, H355 • PERC 12G: H965i, H965e • HBA 11: HBA355i, HBA355e • BOSS-N1 • Software RAID: S160	• PERC 9G: H330, H730p • PERC 10G: H740p, H840 • HBA 9: HBA330, 12 Gbps SAS HBA • HBA 11: HBA355e • BOSS adapter • Software RAID: S140 • Internal Dual SD Module
Drive Bays	Front bays: • 2.5 inches - 24Gb SAS, 6Gb SATA, Gen3/4 NVMe • E3.S - NVMe Gen5 EDSFF	Front bays: • 2.5 inches - 12Gb SAS, 6Gb SATA, Gen3 NVMe
Power Supplies	• AC (Platinum): 1400 W, 2400 W • AC (Titanium): 1100 W, 1800 W, 2800 W	• AC (Platinum): 750 W, 1100 W, 1600 W, 2000 W, 2400 W • AC (Titanium): 1600 W, 2600 W • LVDC @-48VDC Input: 1100 W
Cooling Options	• Air Cooling	• Air Cooling
Fans	Standard (STD) fans	Standard (STD) fans
	Up to 6 sets (dual fan module) hot plug fans	Up to six hot swap fans
Dimension	Height: 174.3 mm (6.86 inches)	Height: 174.3 mm (6.86 inches)
	Width: 482 mm (18.97 inches)	Width: 441.16 mm (17.37-inches)
	Depth: 883.195 mm (34.77 inches) with bezel	Depth: 877.84 mm (34.56 inches) with bezel
	869.195 mm (34.22 inches) without bezel	865.9 (34.09-inches) without bezel
Form Factor	4U rack server	4U rack server
Embedded Management	• iDRAC9 • iDRAC Direct • iDRAC RESTful with Redfish • iDRAC Service Manual • Quick Sync 2 wireless module	• iDRAC9 • iDRAC Direct • Quick Sync 2 BLE/wireless module
Bezel	Optional LCD bezel or security bezel	Optional LCD bezel or security bezel

Table 2. Features comparison (continued)

Features	PowerEdge R960		PowerEdge R940xa	
OpenManage Software	• CloudIQ for PowerEdge plug in • OpenManage Enterprise • OpenManage Enterprise Integration for VMware vCenter • OpenManage Integration for Microsoft System Center • OpenManage Integration with Windows Admin Center • OpenManage Power Manager plugin • OpenManage Service plugin • OpenManage Update Manager plugin		• OpenManage Enterprise • OpenManage Power Center	
Mobility	OpenManage Mobile		OpenManage Mobile	
Integrations and Connections	OpenManage Integrations • BMC TrueSight • Microsoft System Center • OpenManage Integration with ServiceNow • Red Hat Ansible Modules • Terraform Providers • VMware vCenter and vRealize Operations Manager		OpenManage Integrations • BMC TrueSight • Microsoft System Center • Red Hat Ansible Modules • VMware vCenter	• IBM Tivoli Netcool/ OMNIbus • Micro Focus Operations Manager I • Nagios Core • Nagios XI
Security	• Cryptographically signed firmware • Secure Boot • Secure Erase • Silicon Root of Trust • System Lockdown (requires iDRAC9 Enterprise or Datacenter) • TPM 2.0 FIPS, CC-TCG certified, TPM 2.0 China NationZ • Secured Component Verification (Hardware integrity check) • Data at Rest Encryption (SEDs with local or external key mgmt)		• Cryptographically signed firmware • Secure Boot • Secure Erase • Silicon Root of Trust • System Lockdown (requires iDRAC9 Enterprise or Datacenter) • TPM 1.2/2.0 FIPS, CC-TCG certified, TPM 2.0 China NationZ	
Embedded NIC	2 x 1 GbE LOM (optional)		4 x 1 GbE Network Daughter Card (optional)	
Networking Options	OCP x16 (optional) Mezz 3.0 **NOTE:** The system allows either LOM card or an OCP card or both to be installed in the system.		NA	
GPU Options	Up to four double wide 400 W		Up to two double wide 300 W, or eight single wide 75 W accelerators	
Ports	Front Ports • 1 x USB 2.0 • 1 x VGA • 1 x iDRAC Direct (Micro-AB USB) port	Rear Ports • 1 x USB 2.0 • 1 x Dedicated iDRAC Ethernet port • 1 x USB 3.0 • 1 x Serial port (optional) • 1 x VGA	Front Ports • 1 x Dedicated iDRAC micro-USB • 2 x USB 2.0 • 1 x VGA	Rear Ports • 2 x USB 3.0 • 1 x Dedicated iDRAC Ethernet port • 1 x Serial port • 1 x VGA
	Internal Port: 1 x USB 3.0 (optional)		Internal Port: 1 x USB 3.0 (optional)	
PCIe	Up to 12 x PCIe Gen5 slots		Up to 12 x Gen3 slots, (6x16 slots or 2x16 + 10x8 slots)	

Table 2. Features comparison (continued)

Features	PowerEdge R960	PowerEdge R940xa
Operating System and Hypervisors	• Canonical Ubuntu Server LTS • Microsoft Windows Server with Hyper-V • Red Hat Enterprise Linux • SUSE Linux Enterprise Server • VMware ESXi For specifications and interoperability details, see Dell Enterprise Operating Systems on Servers, Storage, and Networking page at Dell.com/OSsupport.	• Canonical Ubuntu Server LTS • Citrix Hypervisor • Windows Server LTSC with Hyper-V • Red Hat Enterprise Linux • SUSE Linux Enterprise Server • VMware ESXi For specifications and interoperability details, see Dell Enterprise Operating Systems on Servers, Storage, and Networking page at Dell.com/OSsupport.

NOTE: OpenManage will be offered for R960 n future release.

3

Chassis views and features

Topics:

- Front view of the system
- Rear view of the system
- Inside the system
- Quick Resource Locator

Front view of the system

Figure 1. Front view of 32 x 2.5-inch drive system

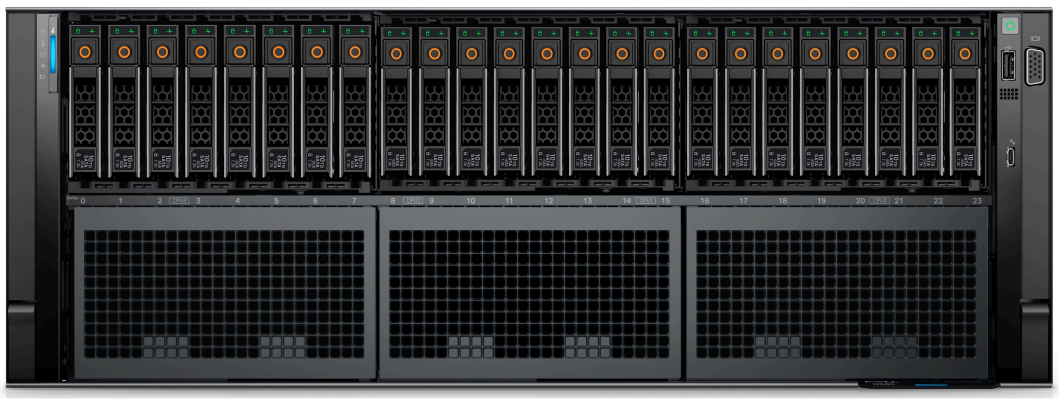

Figure 2. Front view of 24 x 2.5-inch drive system

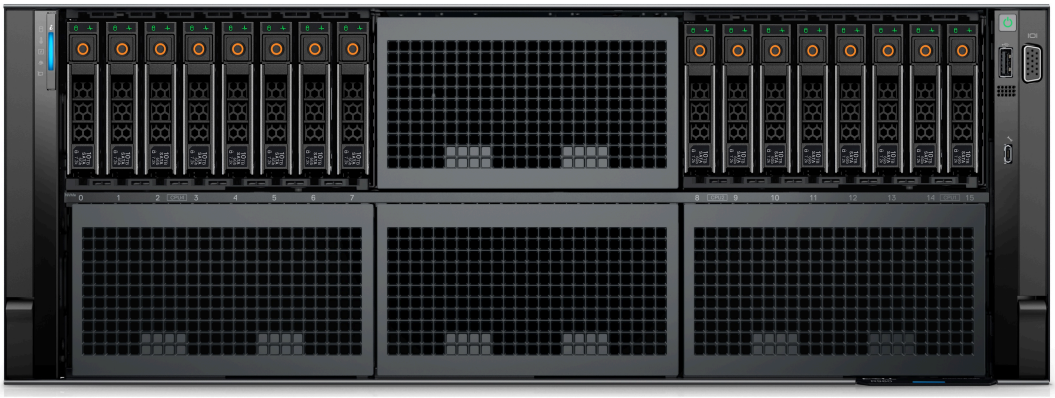

Figure 3. Front view of 16 x 2.5-inch drive system

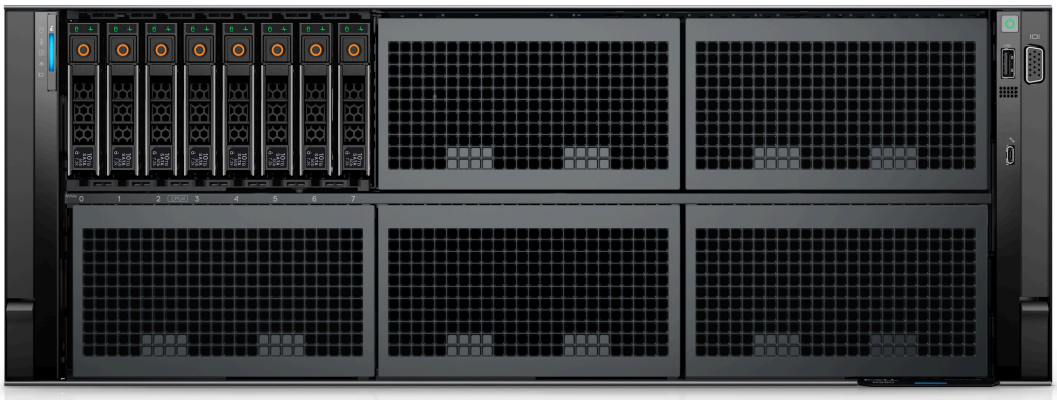

Figure 4. Front view of 8 x 2.5-inch drive system

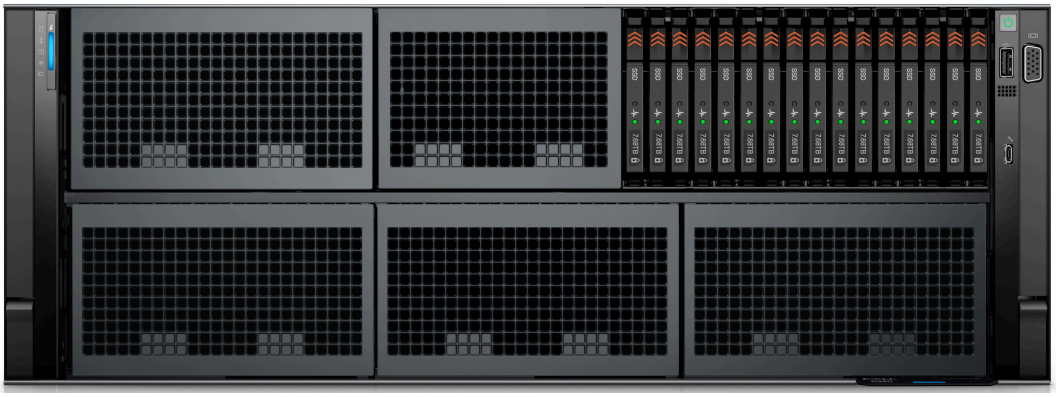

Figure 5. Front view of 16 x EDSFF E3.S Gen5 NVMe drive system

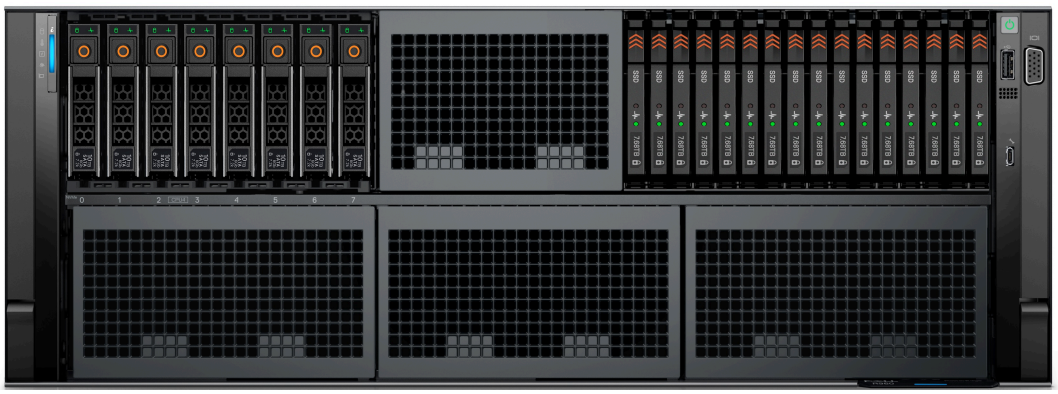

Figure 6. Front view of 16 x EDSFF E3.S Gen5 NVMe with 8 x 2.5-inch drive system

Rear view of the system

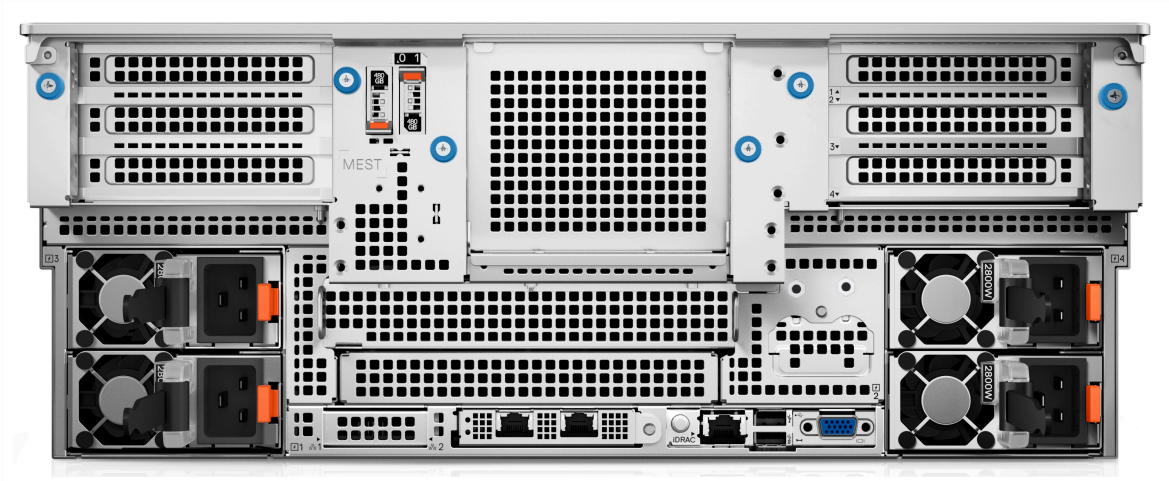

Figure 7. Rear view of the system with basic configuration

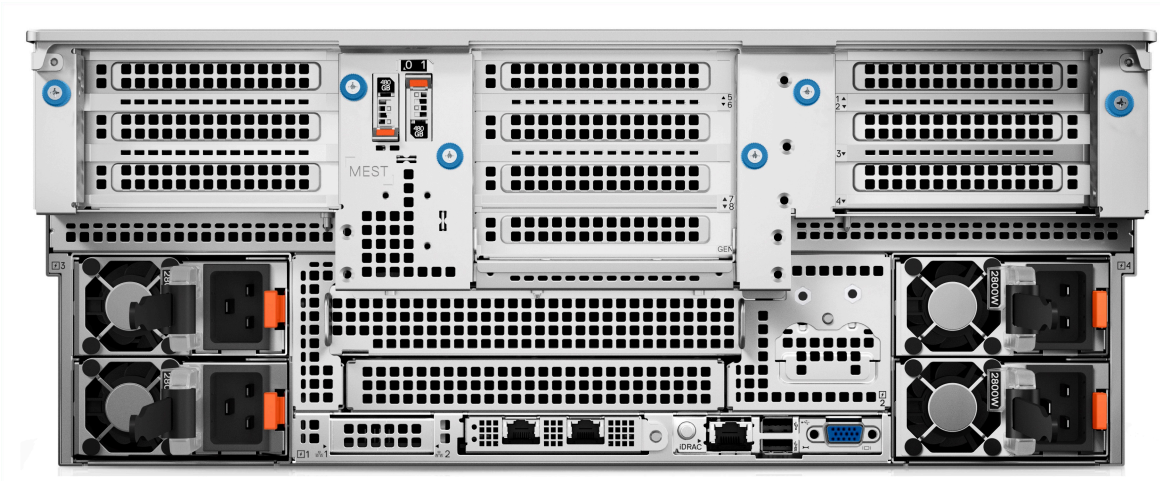

Figure 8. Rear view of the system with scale out configuration

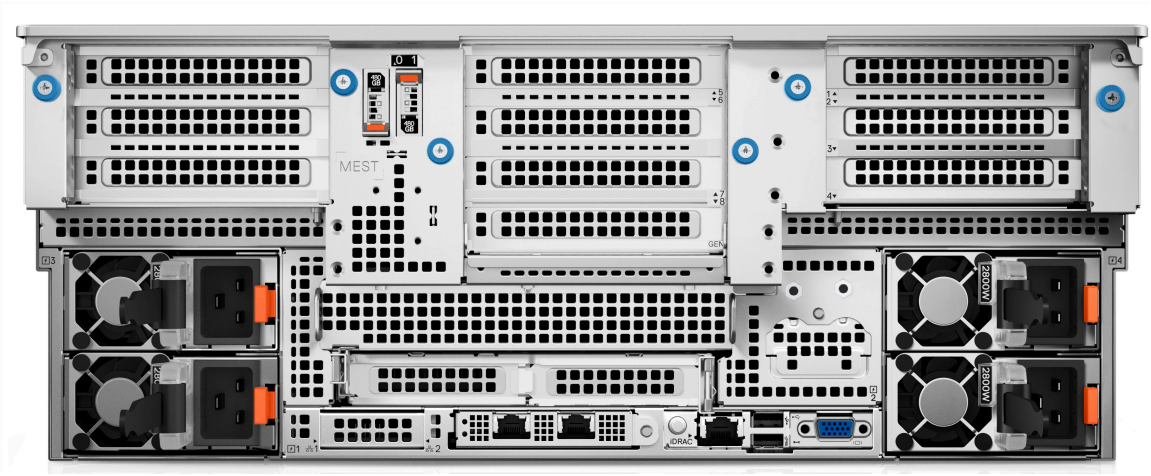

Figure 9. Rear view of the system with maximum I/O configuration

Inside the system

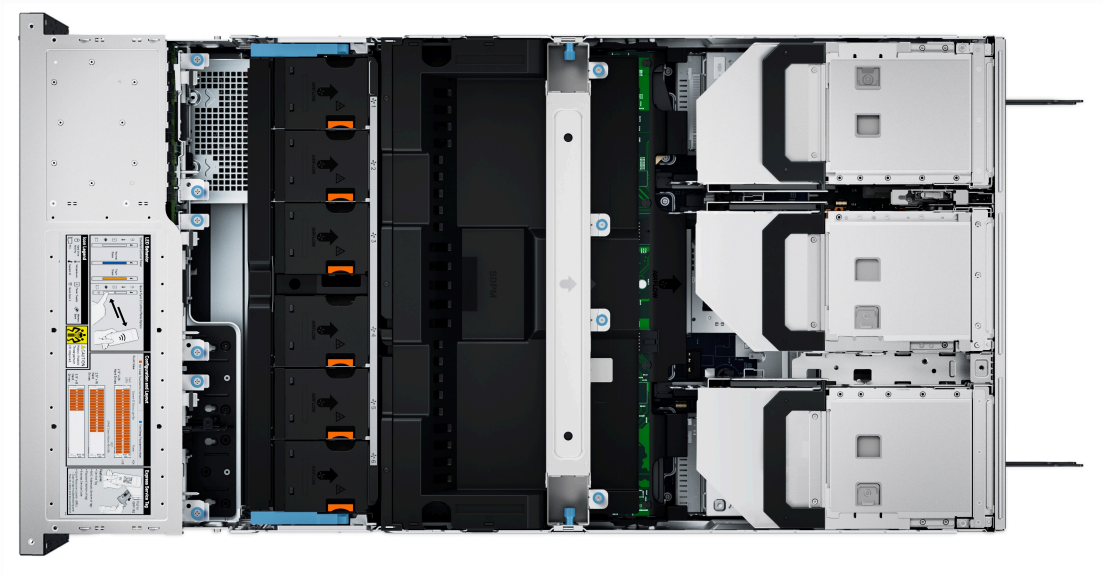

Figure 10. Inside the system

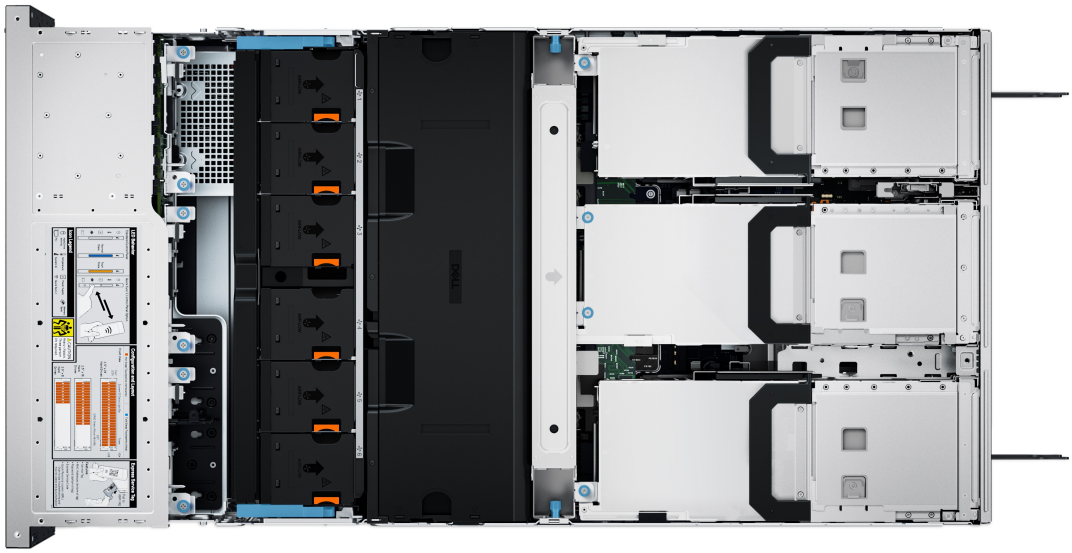

Figure 11. Inside the system with full length risers and GPU shroud

Quick Resource Locator

The QRL on everything (SILs, GSG, Owner's Manual except on the EST) is a generic QRL for R960 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 QRL 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 QRL code within it are printed on demand at the L10 factories. This QRL 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 R960 that is available in the other QRLs.

Figure 12. Quick Resource Locator for PowerEdge R960 system

4

Processor

Topics:

- Processor features

Processor features

The 4th Generation Intel® Xeon® Processors stack is the next generation data center processor offering with significant performance increases, integrated acceleration, and next generation memory and I/O. Sapphire Rapids accelerate customer usages with unique workload optimizations.

The following lists the features and functions that are in the upcoming 4th Generation Intel® Xeon® Scalable Processor offering:

- Faster UPI with up to four Intel Ultra Path Interconnect (Intel UPI) at up to 16 GT/s, increasing multisocket bandwidth
- More, faster I/O with PCI Express 5 and up to 80 lanes (per socket)
- Enhanced Memory Performance with DDR5 support and memory speed up to 4800 MT/s in one DIMM per channel (1DPC) and 4400 MT/s in two DIMM per channel (2DPC)
- New built-in accelerators for data analytics, networking, storage, crypto, and data compression

Supported processors

The following table shows the Intel Sapphire Rapids SKUs that are supported on the R960.

Table 3. Supported Processors for R960

Processor	Clock Speed (GHz)	Cache (M)	UPI (GT/s)	Cores	Threads	Turbo	Memory Speed (MT/s)	Memory Capacity	TDP
8490H	1.9	113	16	60	120	Turbo	4800	6 TB	350 W
8468H	2.1	105	16	48	96	Turbo	4800	6 TB	330 W
8460H	2.2	105	16	40	80	Turbo	4800	6 TB	330 W
8454H	2.1	83	16	32	64	Turbo	4800	6 TB	270 W
8450H	2	75	16	28	56	Turbo	4800	6 TB	250 W
8444H	2.9	45	16	16	32	Turbo	4800	6 TB	270 W
6448H	2.4	60	16	32	64	Turbo	4800	6 TB	250 W
6434H	3.7	23	16	8	16	Turbo	4800	6 TB	195 W
6418H	2.1	60	16	24	48	Turbo	4800	6 TB	185 W
6416H	2.2	45	16	18	36	Turbo	4800	6 TB	165 W

NOTE: Mixing of the processors in a 4S configurations is not recommended.

5

Memory subsystem

The PowerEdge XR8610t and XR8620t supports up to 8 DIMMs, with up to 512 GB of standard memory and speeds of up to 4800 MT/s.

Topics:

- Supported memory

Supported memory

Table 4. Memory technology comparison

Feature	PowerEdge R960 (DDR5)
DIMM type	RDIMM
Transfer speed	4800 MT/s (1DPC), 4400 MT/s (2DPC)
Voltage	1.1 V

Table 5. Supported memory matrix

DIMM type	Rank	Capacity	DIMM rated voltage and speed	Operating Speed	
				1 DIMM per channel (DPC)	2 DIMMs per channel (DPC)
RDIMM	1 R	16 GB	DDR5 (1.1 V), 4800 MT/s	4800 MT/s	4400 MT/s
	2 R	32 GB, 64 GB	DDR5 (1.1 V), 4800 MT/s	4800 MT/s	4400 MT/s
	4 R	128 GB	DDR5 (1.1 V), 4800 MT/s	4800 MT/s	4400 MT/s
	8 R	256 GB	DDR5 (1.1 V), 4800 MT/s	4800 MT/s	4400 MT/s

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

6

Storage

Topics:

- Storage controllers
- Supported Drives
- Internal storage configuration
- External Storage

Storage controllers

Dell RAID controller options offer performance improvements, including the fPERC solution. fPERC provides a base RAID HW controller without consuming a PCIe slot by using a small form factor and high-density connector to the base planar.

16G PERC Controller offerings are a heavy leverage of the 15G PERC family. The Value and Value Performance levels carry over to 16G from 15G. New to 16G is the Avenger-based Premium Performance tier offering. This high-end offering drives IOPs performance and enhanced SSD performance.

Table 6. PERC Series controller offerings

Performance Level	Controller and Description
Entry	S160
Value	H355, HBA355 (internal/external)
Value Performance	H755
Premium Performance	H965i, H965e

NOTE: PowerEdge does not support Tri-Mode, the mixing of SAS, SATA, and NVMe behind the same controller.

NOTE: For more information about the features of the Dell PowerEdge RAID controllers (PERC), Software RAID controllers, or BOSS card, and on deploying the cards, see the storage controller documentation at www.dell.com/storagecontrollermanuals.

NOTE: From December 2021, H355 replaces H345 as the entry raid controller. H345 is deprecated in January 2022.

Supported Drives

The table shown below lists the internal drives supported by the R960.

Table 7. Supported Drives

Form Factor	Type	Speed	Rotational Speed	Capacities
2.5 inches	vSAS	12 Gb	SSD	1.92 TB, 3.84 TB, 960 GB, 7.62 TB
2.5 inches	SAS	24 Gb	SSD	1.92 TB, 1.6 TB, 800 GB, 3.84 TB, 960 GB, 7.68 TB
2.5 inches	SATA	6 Gb	SSD	1.92 TB, 480 GB, 960 GB, 3.84 TB
2.5 inches	NVMe	Gen4	SSD	1.6 TB, 3.2 TB, 6.4 TB, 1.92 TB, 3.84 TB, 15.63 TB, 7.68 TB, 800 GB, 400 GB

Table 7. Supported Drives (continued)

Form Factor	Type	Speed	Rotational Speed	Capacities
2.5 inches	DC NVMe	Gen4	SSD	3.84 TB, 960 GB
2.5 inches	SAS	12 Gb	10 K	600 GB, 1.2 TB, 2.4 TB
EDSFF E3.S	NVMe	Gen5	SSD	3.84 TB, 7.68 TB

Internal storage configuration

R960 available internal storage configurations:

- 8 x 2.5" (SAS/SATA) - Smartflow
- 16 x 2.5" (SAS4/SATA) - Smartflow
- 16 x 2.5"(SAS4/SATA) + 8 x 2.5" NVMe
- 32 x 2.5" (SAS4/SATA) - Dual Controller
- 24 x 2.5" (NVMe Gen4) Passive
- 16 x EDSFF E3.S (NVMe Gen5)
- 16 x EDSFF E3.S (NVMe Gen5) + 8 x 2.5" (SAS4/SATA)

NOTE: The Universal Backplane supports HW RAID for SAS/SATA with direct attach NVMe, and does not support HW RAID for NVMe.

External Storage

The R960 support the external storage device types listed in the table below.

Table 8. Support External Storage Devices

Device Type	Description
External Tape	Supports connection to external USB tape products
NAS/IDM appliance software	Supports NAS software stack
JBOD	Supports connection to 12 Gb MD-series JBODs

7

Networking

Topics:

- Overview
- OCP 3.0 support

Overview

PowerEdge offers a wide variety of options to get information moving to and from our servers. Industry best technologies are chosen, and systems management features are added by our partners to firmware to tie in with iDRAC. These adapters are rigorously validated for worry-free, fully supported use in Dell servers.

OCP 3.0 support

Table 9. OCP 3.0 feature list

Feature	OCP 3.0
Form factor	SFF
PCIe Gen	Gen4
Max PCIe width	x16
Max no. of ports	4
Port type	BT/SFP/SFP+/SFP28/SFP56
Max port speed	25 GbE
NC-SI	Yes
SNAPI	Yes
WoL	Yes
Power consumption	15 W–150 W

Supported OCP cards

Table 10. Supported OCP cards

Form factor	Vendor	Port type	Port speed	Port count
OCP 3.0	Broadcom	SFP28	25 GbE	4
	Broadcom	SFP28	25 GbE	2
	Intel	SFP28	25 GbE	4
	Intel	SFP28	25 GbE	2
	Broadcom	BT	10 GbE	4
	Broadcom	SFP+	10 GbE	2

Table 10. Supported OCP cards (continued)

Form factor	Vendor	Port type	Port speed	Port count
	Broadcom	BT	1 GbE	4
	Intel	SFP	1 GbE	4

OCP NIC 3.0 vs. rack Network Daughter Card comparisons

Table 11. OCP 3.0, 2.0, and rNDC NIC comparison

Form Factor	Dell rNDC	OCP 2.0 (LOM Mezz)	OCP 3.0	Notes
PCIe Gen	Gen 3	Gen 3	Gen 4	Supported OCP3 are SFF (small form factor)
Max PCIe Lanes	x8	Up to x16	Up to x16	See server slot priority matrix
Shared LOM	Yes	Yes	Yes	This is iDRAC port redirect
Aux Power	Yes	Yes	Yes	Used for Shared LOM

8

PCIe subsystem

Topics:

- PCIe risers

PCIe risers

Shown below are the riser offerings for the platform.

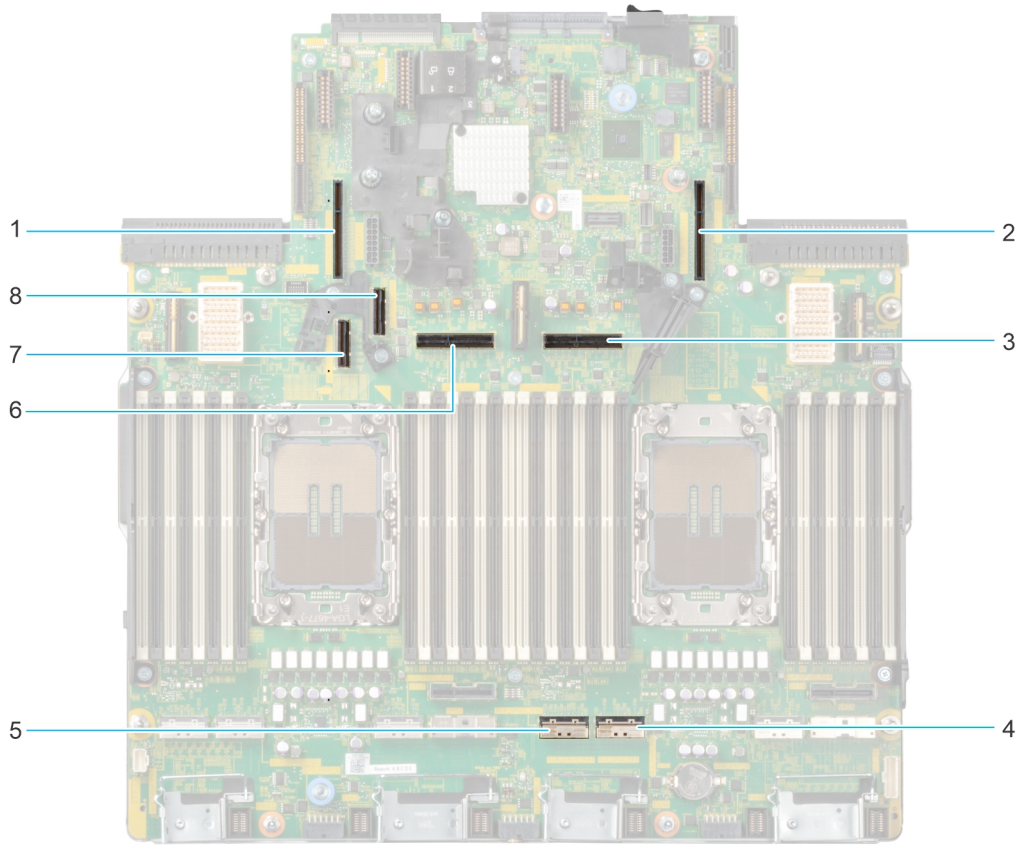

Figure 13. Riser connector location on system board

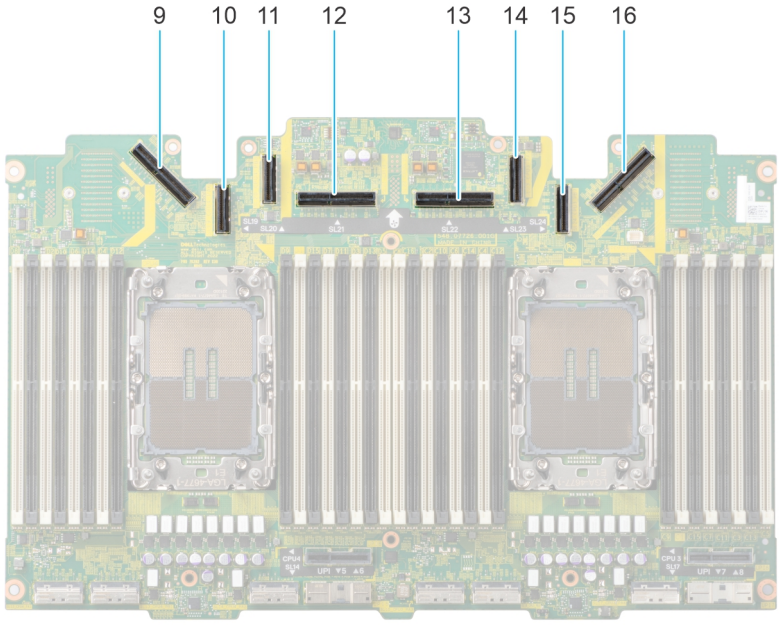

Figure 14. Riser connector location on PEM board

1. Riser 4A/4B/4C slot
2. Riser 1A/1B/1C slot
3. Riser 2 cable slot / Riser 3B cable slot

   NOTE: If present, riser 2 uses the connector, otherwise riser 3B uses it.
4. Riser 1B cable slot
5. Riser 1B cable slot
6. Riser 2 cable slot / Riser 3B cable slot

   NOTE: If present, riser 2 uses the connector, otherwise riser 3B uses it.
7. Riser 4B cable slot
8. Riser 4B cable slot
9. Riser 4A/4B/4C cable slot
10. Riser 3B cable slot
11. Riser 3B cable slot
12. Riser 3A/3B cable slot

    NOTE: If riser 2 is present, riser 3A/3B uses this connector.
13. Riser 3A/3B cable slot

    NOTE: If riser 2 is present, riser 3A/3B uses this connector.
14. Riser 3B cable slot
15. Riser 3B cable slot
16. Riser 1A/1B/1C cable slot

NOTE: A riser can be connected to the riser slot or riser cable slot or both depending on the type of riser used.

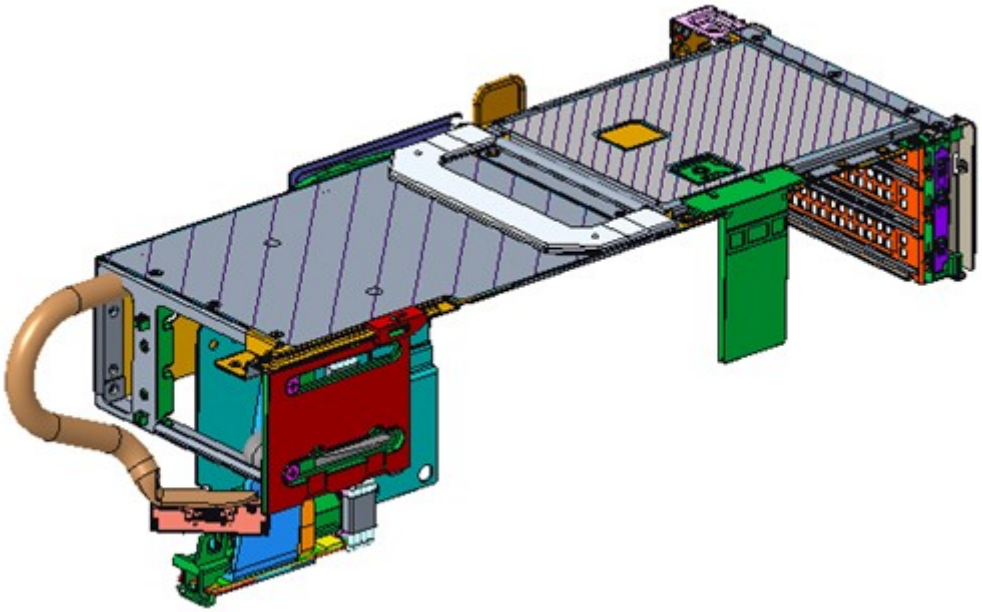

Figure 15. Riser 1A (FL)

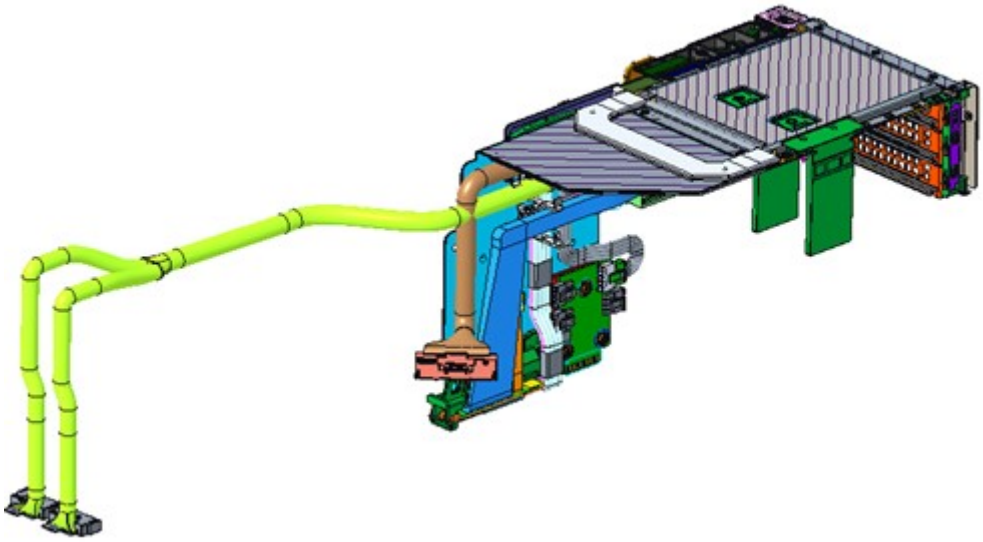

Figure 16. Riser 1B

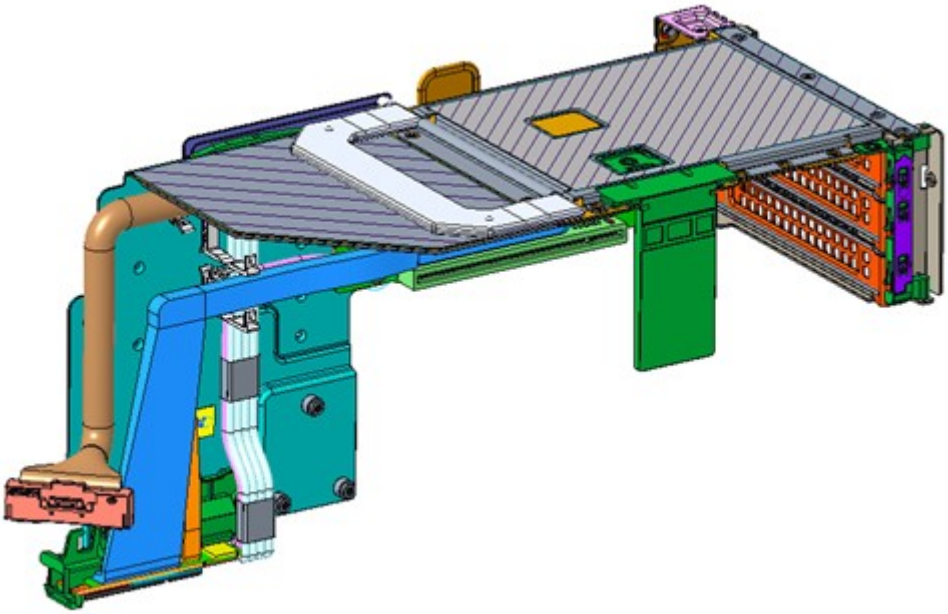

Figure 17. Riser 1C

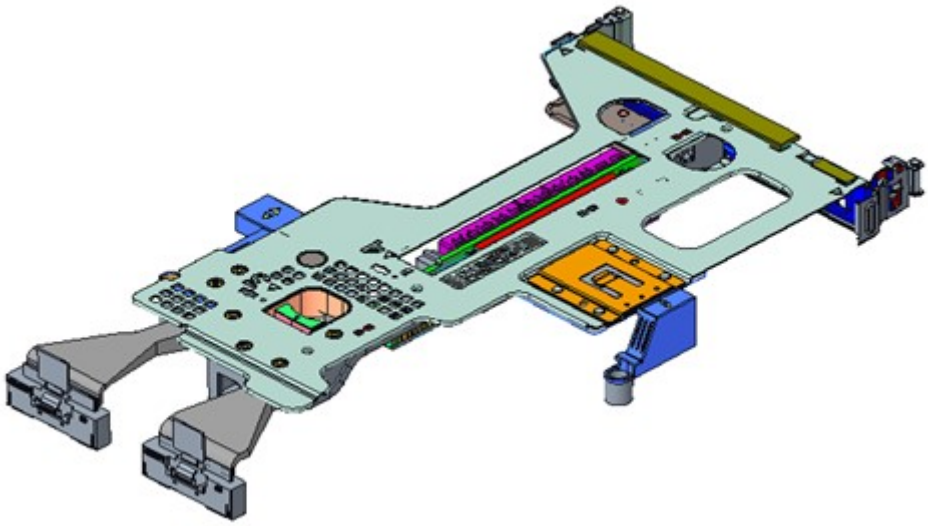

Figure 18. Riser 2

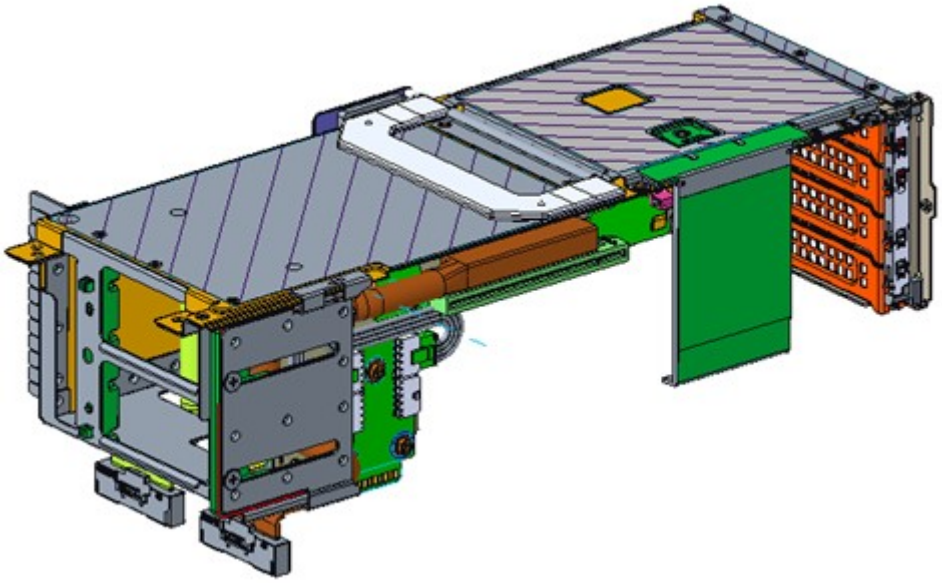

Figure 19. Riser 3A (FL)

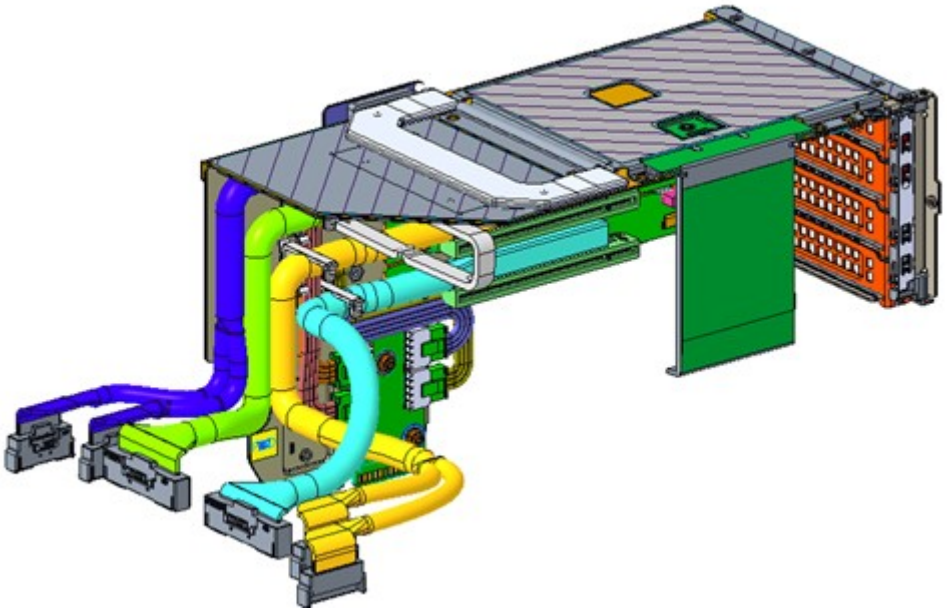

Figure 20. Riser 3B

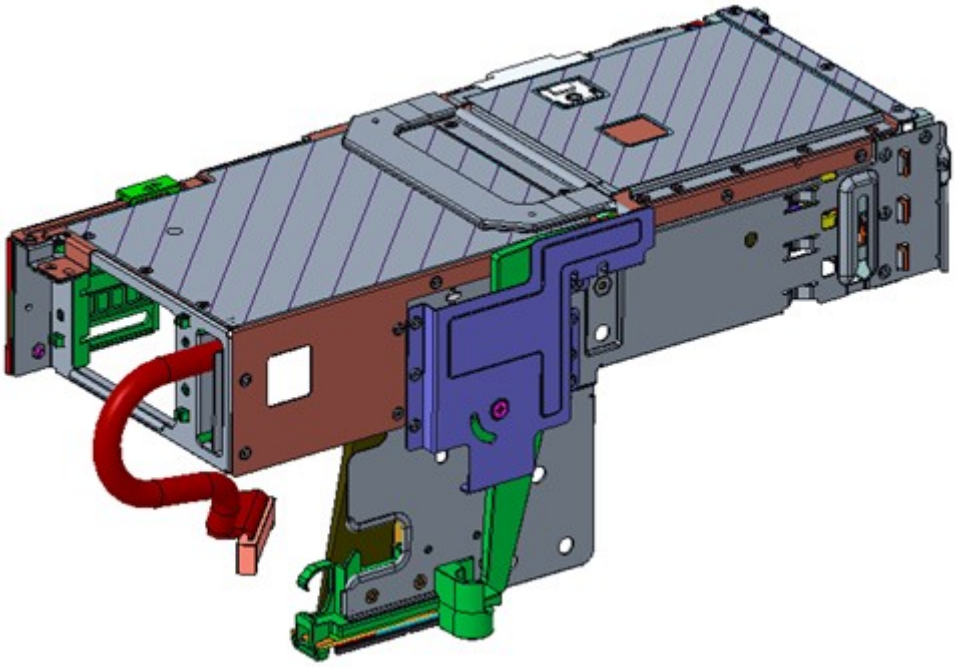

Figure 21. Riser 4A (FL)

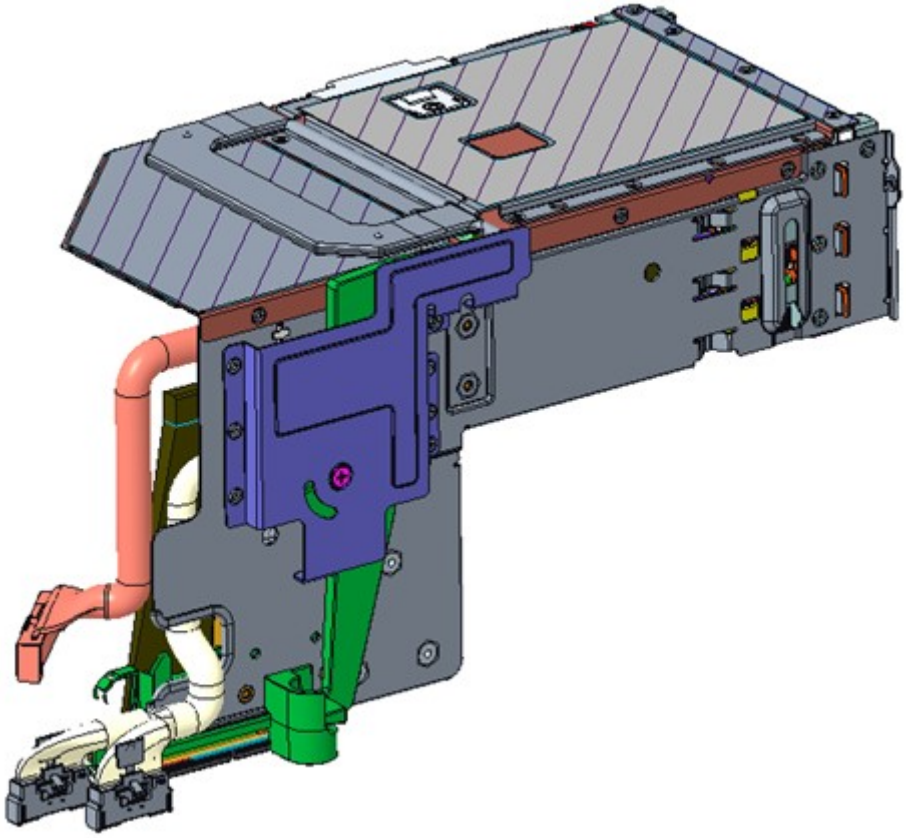

Figure 22. Riser 4B

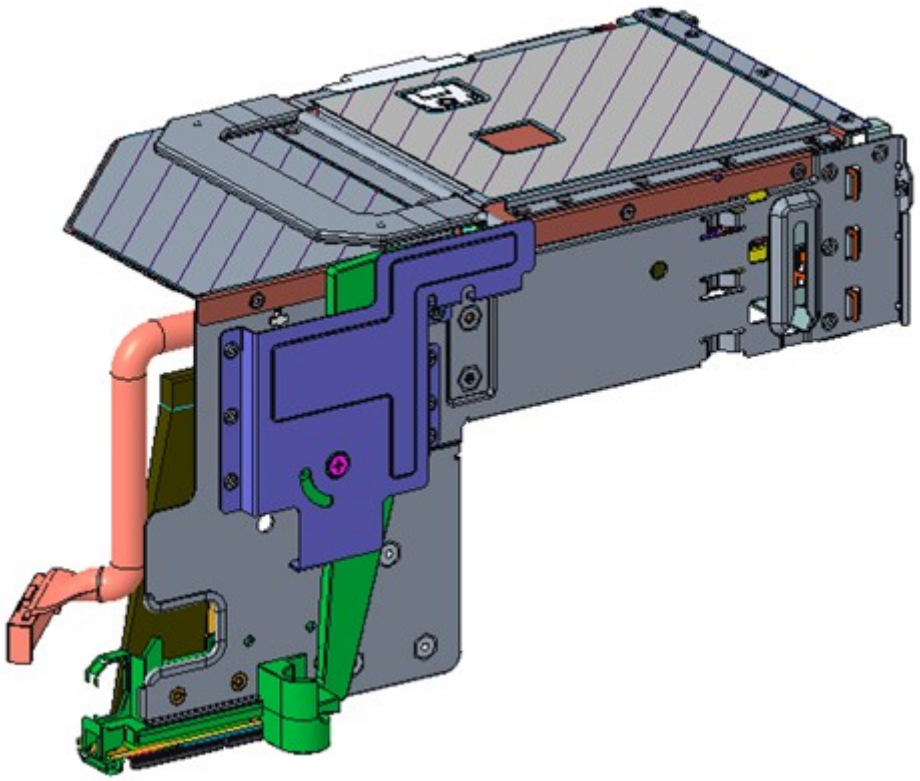

Figure 23. Riser 4C

Table 12. PCIe Riser Configurations

Config No.	Riser configuration	No. of Processors	PERC type supported	Rear storage possible
0-1 (Basic entry)	R1C+R4C	4	Front PERC	No
0-2 (Basic entry)	R1C+R4C	2	Front PERC	No
1 (Scale out)	R1C+R3B+R4C	4	Front PERC	No
2 (Maximum I/O)	R1B+R2+R3B+R4B	4	Front PERC	No
3 (GPU)	R1A+R2+R3A+R4A (FL)	4	Front PERC	No

Table 13. Riser Config 0-1: R1C+R4C

Supporting Processor				CPU3	CPU1	CPU4	CPU2
Location	Width	Length	Height	R1C		R4C	
PCIe Slot-1	SW	HL	FH	Gen5 x16	-	-	-
PCIe Slot-3	SW	HL	FH	-	Gen5 x16	-	-
PCIe Slot-10	SW	HL	FH	-	-	Gen5 x16	-
PCIe Slot-11	SW	HL	FH	-	-	-	Gen5 x16

Table 14. Riser Config 0-2: R1C+R4C

Supporting Processor				CPU3	CPU1	CPU4	CPU2
Location	Width	Length	Height	R1C		R4C	
PCIe Slot-1	SW	HL	FH	NA	-	-	-
PCIe Slot-3	SW	HL	FH	-	Gen5 x16	-	-
PCIe Slot-10	SW	HL	FH	-	-	NA	-

Table 14. Riser Config 0-2: R1C+R4C (continued)

Supporting Processor				CPU3	CPU1	CPU4	CPU2
Location	**Width**	**Length**	**Height**	**R1C**		**R4C**	
PCIe Slot-11	SW	HL	FH	-	-	-	Gen5 x16

Table 15. Riser Config 1: R1C+R3B+R4C

Supporting Processor				CPU3	CPU1	CPU4	CPU2	CPU3	CPU1	CPU4	CPU2
Location	**Width**	**Length**	**Height**	**R1C**		**R3B**				**R4C**	
PCIe Slot-1	SW	HL	FH	Gen5 x16	-	-	-	-	-	-	-
PCIe Slot-3	SW	HL	FH	-	Gen5 x16	-	-	-	-	-	-
PCIe Slot-5	SW	HL	FH	-	-	Gen5 x16	-	-	-	-	-
PCIe Slot-6	SW	HL	FH	-	-	-	Gen5 x16	-	-	-	-
PCIe Slot-7	SW	HL	FH	-	-	-	-	Gen5 x16	-	-	-
PCIe Slot-8	SW	HL	FH	-	-	-	-	-	Gen5 x16	-	-
PCIe Slot-10	SW	HL	FH	-	-	-	-	-	-	Gen5 x16	-
PCIe Slot-11	SW	HL	FH	-	-	-	-	-	-	-	Gen5 x16

Table 16. Riser Config 2: R1B+R2+R3B+R4B

Supporting Processor				CPU1	CPU3	CPU1	CPU2	CPU3	CPU4		CPU2
Location	**Width**	**Length**	**Height**	**R1B**		**R2**		**R3B**		**R4B**	
PCIe Slot-1	SW	HL	FH	-	Gen5 x16	-	-	-	-	-	-
PCIe Slot-2	SW	HL	FH	Gen5 x16	-	-	-	-	-	-	-
PCIe Slot-3	SW	HL	FH	Gen5 x16	-	-	-	-	-	-	-
PCIe Slot-4	SW	HL	LP	-	-	-	Gen5 x16	-	-	-	-
PCIe Slot-5	SW	HL	FH	-	-	-	-	-	Gen5 x16	-	-
PCIe Slot-6	SW	HL	FH	-	-	-	-	-	Gen5 x16	-	-
PCIe Slot-7	SW	HL	FH	-	-	-	-	Gen5 x16	-	-	-
PCIe Slot-8	SW	HL	FH	-	-	-	-	Gen5 x16	-	-	-
PCIe Slot-9	SW	HL	LP	-	-	Gen5 x16	-	-	-	-	-
PCIe Slot-10	SW	HL	FH	-	-	-	-	-	-	Gen5 x16	-
PCIe Slot-11	SW	HL	FH	-	-	-	-	-	-	-	Gen5 x16
PCIe Slot-12	SW	HL	FH	-	-	-	-	-	-	-	Gen5 x16

Table 17. Riser Config 3: R1A+R2+R3A+R4A

Supporting Processor				CPU1	CPU3	CPU1	CPU2	CPU3	CPU4		CPU2
Location	**Width**	**Length**	**Height**	**R1A**		**R2**		**R3A**		**R4A**	
PCIe Slot-1	SW	FL	FH	-	Gen5 x16	-	-	-	-	-	-
PCIe Slot-3	DW	FL	FH	Gen5 x16	-	-	-	-	-	-	-
PCIe Slot-4	SW	HL	LP	-	-	-	Gen5 x16	-	-	-	-
PCIe Slot-6	DW	FL	FH	-	-	-	-	-	Gen5 x16	-	-
PCIe Slot-8	DW	FL	FH	-	-	-	-	Gen5 x16	-	-	-
PCIe Slot-9	SW	HL	LP	-	-	Gen5 x16	-	-	-	-	-
PCIe Slot-10	SW	FL	FH	-	-	-	-	-	-	Gen5 x16	-
PCIe Slot-11	DW	FL	FH	-	-	-	-	-	-	-	Gen5 x16

9

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:

- Power
- Thermal
- Acoustics

Power

Table 18. Power tools and technologies

Feature	Description
Power Supply Units(PSU) portfolio	Dell's PSU portfolio includes intelligent features such as dynamically optimizing efficiency while maintaining availability and redundancy. Find additional information in the Power supply units section.
Tools for right sizing	Enterprise Infrastructure Planning Tool (EIPT) is a tool that can help you determine the most efficient configuration possible. With Dell's EIPT, you can calculate the power consumption of your hardware, power infrastructure, and storage at a given workload. Learn more at www.dell.com/calc.
Industry Compliance	Dell's servers are compliant with all relevant industry certifications and guide lines, including 80 PLUS, Climate Savers and ENERGY STAR.
Power monitoring accuracy	PSU power monitoring improvements include: ● Dell's power monitoring accuracy is currently 1%, whereas the industry standard is 5% ● More accurate reporting of power ● Better performance under a power cap
Power capping	Use Dell's systems management to set the power cap limit for your systems to limit the output of a PSU and reduce system power consumption. Dell is the first hardware vendor to leverage Intel Node Manager for circuit-breaker fast capping.
Systems Management	iDRAC Enterprise and Datacenter provides server-level management that monitors, reports and controls power consumption at the processor, memory and system level. Dell OpenManage Power Center delivers group power management at the rack, row, and data center level for servers, power distribution units, and uninterruptible power supplies.
Active power management	Intel Node Manager is an embedded technology that provides individual server-level power reporting and power limiting functionality. Dell offers a complete power management solution comprised of Intel Node Manager accessed through Dell iDRAC9 Datacenter and OpenManage Power Center that allows policy-based management of power and thermal at the individual server, rack, and data center level. Hot spare reduces power consumption of redundant power supplies. Thermal control off a speed optimizes the thermal settings for your environment to reduce fan consumption and lower system power consumption. Idle power enables Dell servers to run as efficiently when idle as when at full workload.
Rack infrastructure	Dell offers some of the industry's highest-efficiency power infrastructure solutions, including:

Table 18. Power tools and technologies (continued)

Feature	Description
	• Power distribution units (PDUs) • Uninterruptible power supplies (UPSs) • Energy Smart containment rack enclosures Find additional information at: https://www.delltechnologies.com/en-us/servers/power-and-cooling.htm.

Power Supply Units

Energy Smart power supplies have intelligent features, such as the ability to dynamically optimize efficiency while maintaining availability and redundancy. Also featured are enhanced power-consumption reduction technologies, such as high-efficiency power conversion and advanced thermal-management techniques, and embedded power-management features, including high-accuracy power monitoring. The table below shows the power supply unit options that are available for the R960.

Table 19. Power Supply Unit Options

Wattage	Frequency	Voltage/Current	Class	Heat dissipation
1100 W mixed mode	50/60 Hz	100–240 Vac/12—3.6 A	Titanium	4100 BTU/hr
	N/A	240 Vdc/5.2 A	N/A	4100 BTU/hr
1400 W mixed mode	50/60 Hz	100–240 Vac/12—8 A	Platinum	5250 BTU/hr
	N/A	240 Vdc/6.6 A	N/A	5250 BTU/hr
1800 W mixed mode	50/60 Hz	200–240 Vac/10 A	Titanium	6750 BTU/hr
	N/A	240 Vdc/8.2 A	N/A	6750 BTU/hr
2400 W mixed mode	50/60 Hz	100–240 Vac/ 16—13.5 A	Platinum	9000 BTU/hr
	N/A	240 Vdc/11.2 A	N/A	9000 BTU/hr
2800 W mixed mode	50/60 Hz	200–240 Vac/15.6 A	Titanium	10,500 BTU/hr
	N/A	240 Vdc/13.6 A	N/A	10,500 BTU/hr

NOTE: If a system with AC 2400 W PSUs operates at low line 100-120 Vac, and then the power rating per PSU is degraded to 1400 W.

NOTE: If a system with AC 1400 W or 1100 W PSUs operates at low line 100-120 Vac, and then the power rating per PSU is degraded to 1050 W.

Figure 24. PSU power cords

Table 20. PSU power cords

Form factor	Output	Power cord
Redundant 60 mm	1100 W AC	C13
	1400 W AC	C13
	1800 W AC	C15

Table 20. PSU power cords (continued)

Form factor	Output	Power cord
Redundant 86 mm	2400 W AC	C19
	2800 W AC	C21

NOTE: C19 power cord combined with C20 to C21 jumper power cord can be used to adapt 2800 W PSU.

NOTE: C13 power cord combined with C14 to C15 jumper power cord can be used to adapt 1800 W PSU.

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.

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.

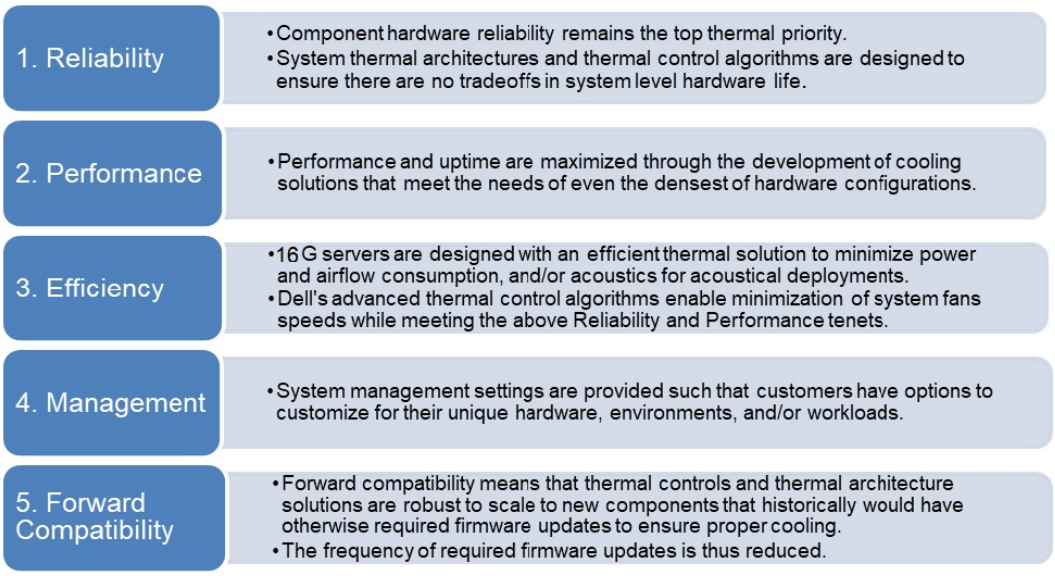

Figure 25. Thermal design characteristics

The thermal design of the PowerEdge R960 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, in this generation of servers, we have introduced limited user- configurable settings residing in the iDRAC BIOS setup screen. For more information, see the Dell PowerEdge R960 Installation and Service Manual at

www.dell.com/poweredgemanuals and “Advanced Thermal Control: Optimizing across Environments and Power Goals” on Dell.com.

- Cooling redundancy: The R960 allows N+1 fan redundancy, allowing continuous operation with one fan failure in the system.
- Environmental Specifications: The optimized thermal management makes the R960 reliable under a wide range of operating environments.

Acoustics

Acoustical configurations of R960

Dell PowerEdge R960 is a rack or tower server appropriate for attended data center environment. However, lower acoustical output is attainable with proper hardware or software configurations.

Table 21. Configurations tested for acoustical experience

Configuration	Quietest	HCI	Data Management	Machine Learning
CPU TDP	Intel (Gold 6) 195 W	Intel (Gold 6) 250 W	Intel (Gold 6) 250 W	Intel (Platinum 8) 350 W
CPU Quantity	4	4	4	4
RDIMM Memory	16G DDR5 4800 MHz	64G DDR5 4800 MHz	64G DDR5 4800 MHz	64G DDR5 4800 MHz
Memory Quantity	4	32	48	64
Backplane Type	2.5" x 32 BP	2.5" x 32 BP	2.5" x 16 BP	2.5" x 32 BP
HDD Type	SAS TOSHIBA AL15SE 2.5" 600 GB	SAS TOSHIBA AL15SE 2.5" 600 GB	SAS TOSHIBA AL15SE 2.5" 1.2 TB	SAS Seagate Skybolt 2.5" 1.8 TB
HDD Quantity	32	32	16	32
M.2 Drives	Hynix NVMe 480G	Hynix NVMe 480G	Hynix NVMe 480G	Hynix NVMe 480G
M.2 Quantity	1	1	1	1
PSU Type	ARTESYN 2400 W	LITEON 2800 W	ARTESYN 2400 W	LITEON 2800 W
PSU Quantity	4	4	4	4
OCP	X	Intel OCP3 10G 2-port	Intel OCP3 10G 2-port	Intel OCP3 10G 2-port
PCI 1	X	X	X	X
PCI 2	X	X	Nvidia A16 250W	X
PCI 3	X	X	X	X
PCI 4	X	X	X	X
PCI 5	X	X	X	X
PCI 6	X	X	X	X
PCI 7	X	X	X	X
PCI 8	X	X	X	X
PCI 9	X	X	X	X
PCI 10	X	X	Nvidia A16 250 W	X
PCI 11	X	X	X	X
PCI 12	X	X	X	X
PERC	Front H755	Front H755	Front H755	Front H755

Table 22. Acoustical experience of R960 configurations

Configuration		Quietest	HCI	Data Management	Machine Learning
Acoustical Performance: Idle/ Operating @ 25°C Ambient					
$L_{wA,m}$ (B)	Idle[4]	5.4	6.0	6.0	5.6
	Operating/Customer usage operating[5][6]	5.7	6.3	7.5	6.7
K_v (B)	Idle [4]	0.4	0.4	0.4	0.4
	Operating/Customer usage operating[5][6]	0.4	0.4	0.4	0.4
$L_{pA,m}$ (dB)	Idle[4]	39	44	44	41
	Operating/Customer usage operating[5][6]	43	47	60	53
Prominent discrete tones[3]		Prominence ratio < 15 dB			
Acoustical Performance: Idle @ 28°C Ambient					
$L_{wA,m}$[1] (B)		6.1	6.6	6.9	6.4
K_v (B)		0.4	0.4	0.4	0.4
$L_{pA,m}$[2] (dB)		45	51	54	49
Acoustical Performance: Max. loading @ 35°C Ambient					
$L_{wA,m}$[1] (B)		7.5	7.8	9.0	9.2
K_v (B)		0.4	0.4	0.4	0.4
$L_{pA,m}$[2](dB)		60	63	74	77

[1]LwA,m: The declared mean A-weighted sound power level (LwA) is calculated per section 5.2 of ISO 9296 (2017) with data collected using the methods that are described in ISO 7779 (2010). Engineering data presented here may not be fully compliant with ISO 7779 declaration requirement.

[2]LpA,m: The declared mean A-weighted emission sound pressure level is at the bystander position per section 5.3 of ISO 9296 (2017) and measured using methods that are described in ISO 7779 (2010). The system is placed in a 24U rack enclosure, 25 cm above a reflective floor. Engineering data presented here may not be fully compliant with ISO 7779 declaration requirement.

[3]Prominent tones: Criteria of Annex D of ECMA-74 and Prominence Ratio method of ECMA-418 are followed to determine if discrete tones are prominent and to report them, if so.

[4]Idle mode: The steady-state condition in which the server is energized but not operating any intended function.

[5]Operating mode: The maximum of the steady state acoustical output at 50% of CPU TDP or active storage drives for the respective sections of Annex C of ECMA-74.

[6] Customer Usage Operating mode: The operating mode is represented by the maximum of the steady state acoustical output at 25%~30% of CPU TDP, 2.5%~10% IOPs load, and >80% GPU load as the components showed in the above configurations.

PowerEdge acoustical specifications

For more information about acoustical specifications, see ENG0019663. (See the category definitions.)

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 23. 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.			

Table 23. Dell Enterprise Category 1, "Table-top 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
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 24. 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			N/A

Table 24. 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**
		○ 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"			
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 and program configuration document	Report for all mics	Report for all mics	Report for all mics	Report for all mics

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 25. 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"
		• 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 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 26. 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.				

Table 26. Dell Enterprise Category 4, "Attended 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 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	
		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 27. 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

Table 27. 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	
	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 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

10

Rack, rails, and cable management

Topics:

- Rails and cable management information

Rails and cable management information

The rail offerings for the PowerEdge R960 consist of one sliding type. The cable management offerings consist of an optional cable management arm (CMA) and an optional strain relief bar (SRB).

See the *Enterprise Systems Rail Sizing and Rack Compatibility Matrix* available at https://i.dell.com/sites/csdocuments/Business_solutions_engineering-Docs_Documents/en/rail-rack-matrix.pdf for information regarding:

- Specific details about rail types.
- Rail adjustability ranges for various rack mounting flange types.
- Rail depth with and without cable management accessories.
- Rack types that are supported for various rack mounting flange types.

Key factors governing proper rail selection include the following:

- Spacing between the front and rear mounting flanges of the rack.
- Type and location of any equipment that is mounted in the back of the rack such as power distribution units (PDUs).
- Overall depth of the rack.

Sliding rails features summary

The sliding rails allow the system to be fully extended out of the rack for service. There is one type of sliding rails available, Stab-in/Drop-in sliding rails. The sliding rails are available with or without the optional cable management arm (CMA) or strain relief bar (SRB).

B19 Stab-in/Drop-in sliding rails for 4-post racks

- Supports drop-in or stab-in installation of the chassis to the rails.
- Support for tool-less installation in 19" EIA-310-E compliant square, unthreaded round hole racks including all generations of the Dell racks. Also supports tool-less installation in threaded round hole 4-post racks.
- Support for tool-less installation in Dell Titan or Titan-D racks.
- Support full extension of the system out of the rack to allow serviceability of key internal components.
- Support for optional cable management arm (CMA).
- Support for optional strain relief bar (SRB).

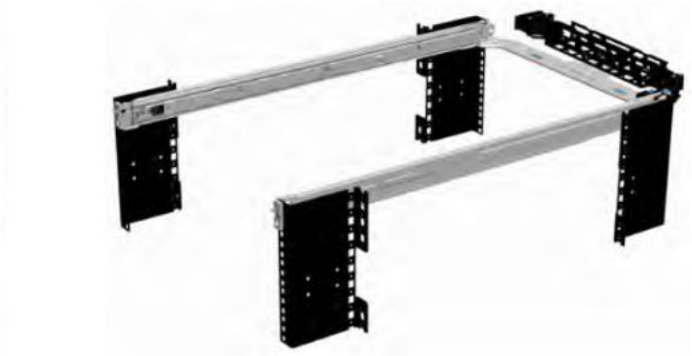

Figure 26. Sliding rails with optional CMA

Figure 27. Sliding rails with optional SRB

Scan the QRL code for the documentation and trouble-shooting information regarding the installation procedures for Drop-in/ Stab-in rail types.

Figure 28. Quick resource locator for combo rails

Installation in the Dell Titan or Titan-D racks

For tool-less installation in Titan or Titan-D racks, the Stab-in/Drop-in sliding rails (B19) must be used. This rail collapses down sufficiently to fit in the rack with mounting flanges that are spaced about 24 inches apart from front to back. The Stab-in/Drop-in sliding rail allows bezels of the servers and storage systems to be aligned when installed in these racks.

Cable management arm (CMA)

The optional cable management arm (CMA) organizes and secures the cords and cables exiting the back of the systems. It unfolds to allow the systems to extend out of the rack without having to detach the cables. Some key features of the CMA include:

- Large U-shaped baskets to support dense cable loads.
- Open vent pattern for optimal airflow.
- Ability to mount on either side by swinging the spring-loaded brackets from one side to the other.
- Utilizes hook-and-loop straps rather than plastic tie wraps to eliminate the risk of cable damage during cycling.
- Includes a low-profile fixed tray to both support and retain the CMA in its fully closed position.
- Both the CMA and the tray mount without the use of tools by simple and intuitive snap-in designs.

For systems with one power supply unit (PSU), it is recommended to mount on the side opposite to that of the power supply to allow easier access to it and the rear drives (if applicable) for service or replacement.

Figure 29. Sliding rails with CMA

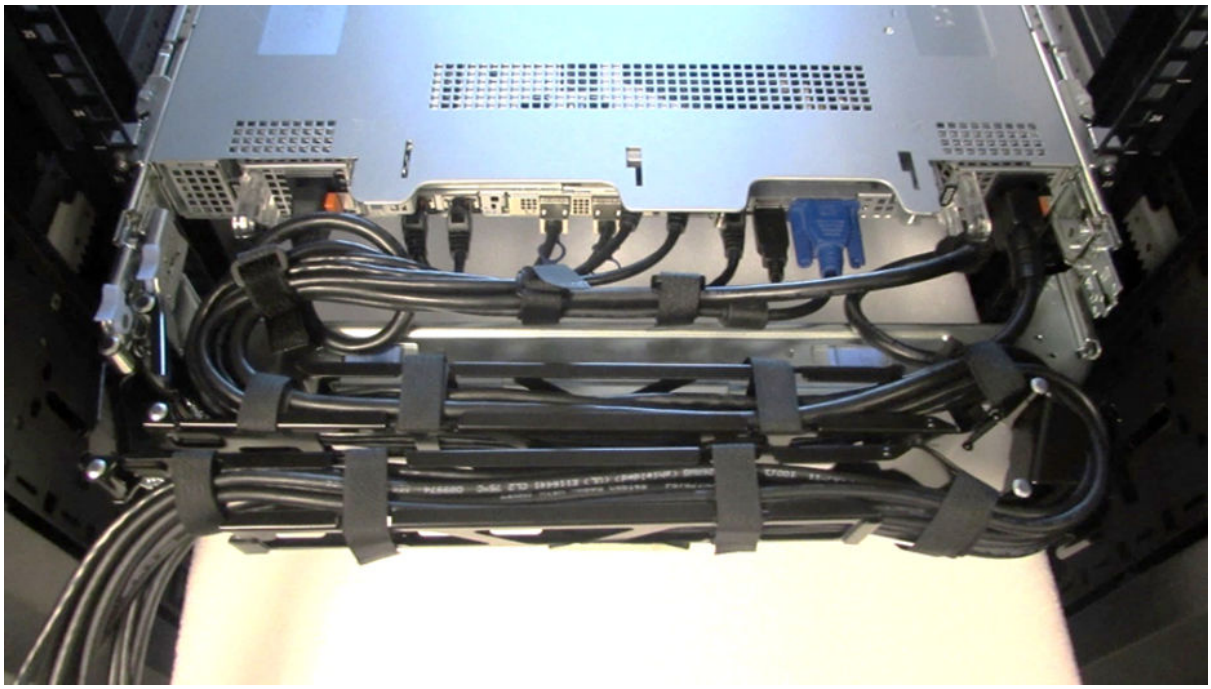

Figure 30. CMA Cabling

Strain Relief Bar (SRB)

The optional strain relief bar (SRB) for the PowerEdge R960 organizes and supports cable connections at the rear end of the server to avoid damage from bending.

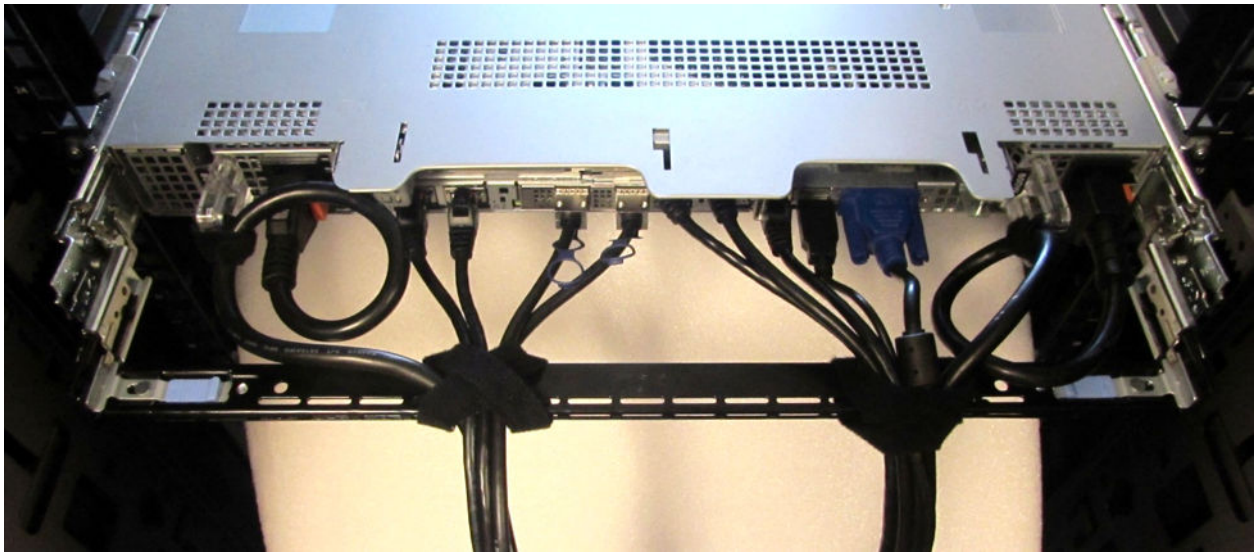

Figure 31. Cabled strain relief bar

- Tool-less attachment to the rails
- Two depth positions to accommodate various cable loads and rack depths
- Supports cable loads and controls stresses on server connectors
- Cables can be segregated into discrete purpose-specific bundles

Rack Installation

Drop-in design means that the system is installed vertically into the rails by inserting the standoffs on the sides of the system into the J-slots in the inner rail members with the rails in the fully extended position. The recommended method of installation is to first insert the rear standoffs on the system into the rear J-slots on the rails to free up a hand and then rotate the system down into the remaining J-slots while using the free hand to hold the rail against the side of the system.

Stab-in design means that the inner (chassis) rail members must first be attached to the sides of the system and then inserted into the outer (cabinet) members installed in the rack.

Installing system into the rack (option A: Drop-In)

1. Pull the inner rails out of the rack until they lock into place.

Figure 32. Pull out inner rail

2. Locate the rear rail standoff on each side of the system and lower them into the rear J-slots on the slide assemblies.
3. Rotate the system downward until all the rail standoffs are seated in the J-slots.

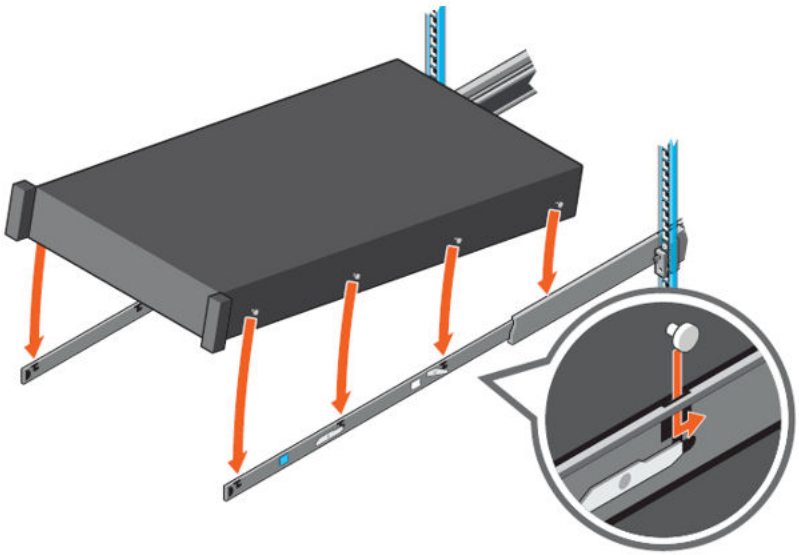

Figure 33. Rail standoffs seated in J-slots

4. Push the system inward until the lock levers click into place.
5. Pull the blue side release lock tabs forward or backward on both rails and slide the system into the rack until the system is in the rack.

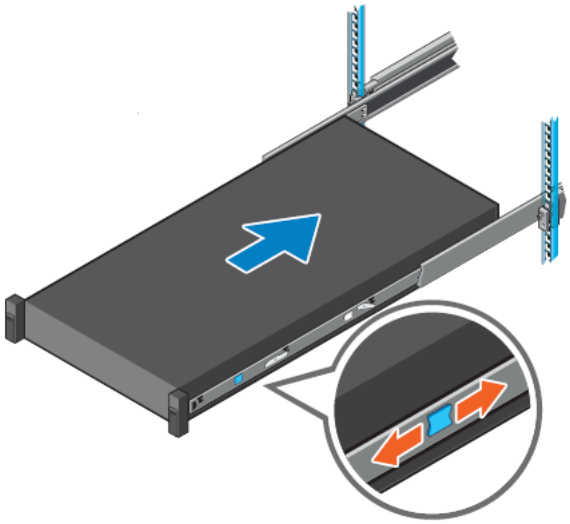

Figure 34. Slide system into the rack

Installing the system into the rack (option B: Stab-In)

1. Pull the intermediate rails out of the rack until they lock into place.
2. Release the inner rail lock by pulling forward on the white tabs and sliding the inner rail out of the intermediate rails.

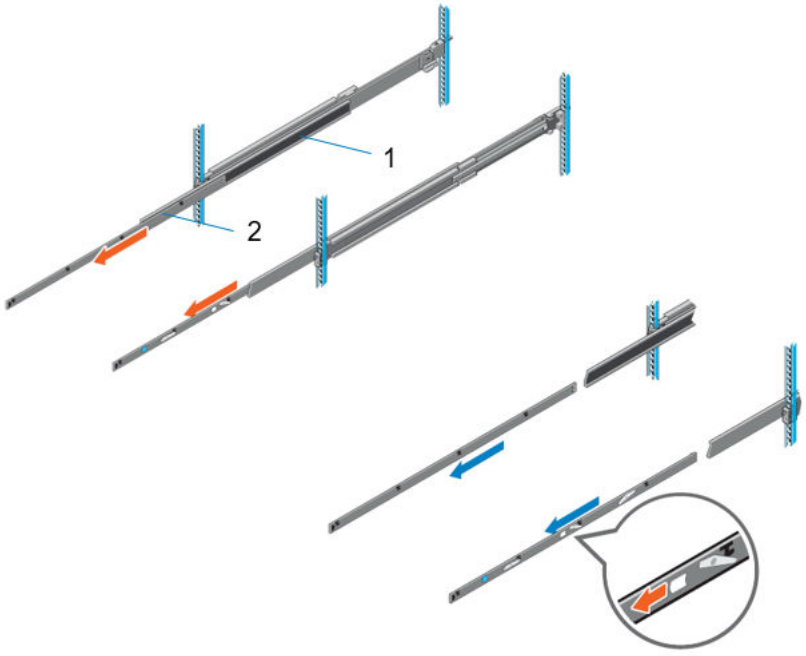

Figure 35. Pull out the intermediate rail

Table 28. Rail component label

Number	Component
1	Intermediate rail
2	Inner rail

3. Attach the inner rails to the sides of the system by aligning the J-slots on the rail with the standoffs on the system and sliding forward on the system until they lock into place.

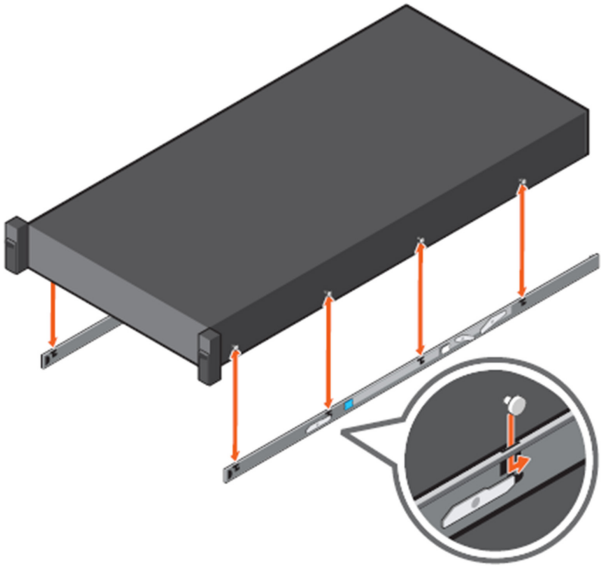

Figure 36. Attach the inner rails to the system

4. With the intermediate rails extended, install the system into the extended rails.

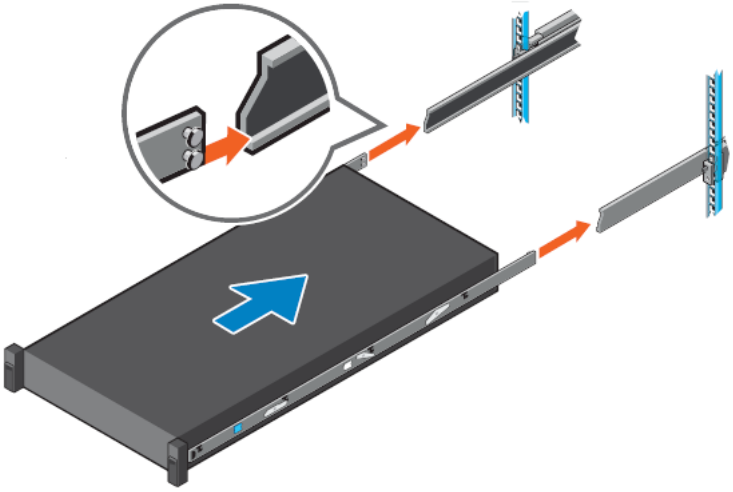

Figure 37. Install system into the extended rails

5. Pull blue slide release lock tabs forward or backward on both rails, and slide the system into the rack.

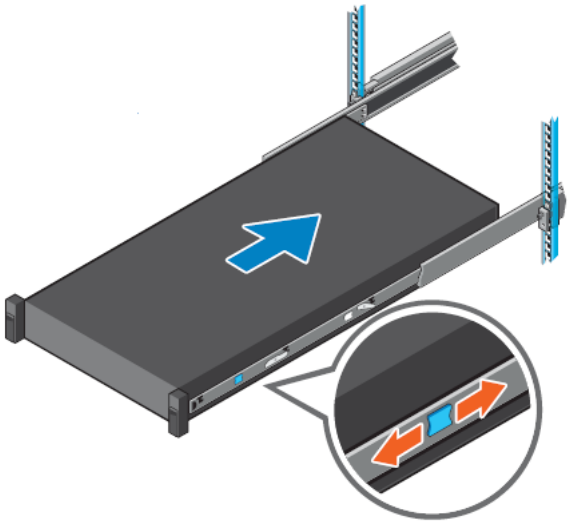

Figure 38. Slide system into the rack

11

Operating Systems and Virtualization

Topics:

- Supported Operating Systems

Supported Operating Systems

The PowerEdge system supports the following operating systems:

- Canonical® Ubuntu® Server LTS
- Microsoft® Windows Server® with Hyper-V
- Red Hat® Enterprise Linux
- SUSE® Linux Enterprise server
- VMware® ESXi®

Links to specific OS versions and editions, certification matrices, Hardware Compatibility Lists (HCL) portal, and Hypervisor support are available at Dell Enterprise Operating Systems.

12

Dell OpenManage Systems Management

Dell delivers management solutions that help IT administrators effectively deploy, update, monitor, and manage IT assets. OpenManage solutions and tools enable you to quickly respond to problems by helping them to manage Dell servers efficiently; in physical, virtual, local, and remote environments; all without the need to install an agent in the operating system.

The OpenManage portfolio includes:

- Innovative embedded management tools - integrated Dell Remote Access Controller (iDRAC)
- Consoles - OpenManage Enterprise
- Extensible with plug-ins - OpenManage Power Manager
- Update tools - Repository Manager

Dell has developed comprehensive systems management solutions that are based on open standards and has integrated with management consoles from partners such as Microsoft and VMware, allowing advanced management of Dell servers. Dell management capabilities extend to offerings from the industry's top systems management vendors and frameworks such as Ansible, Splunk, and ServiceNow. OpenManage tools automate the full span of server life cycle management activities along with powerful RESTful APIs to script or integrate with your choice of frameworks.

For more information about the entire OpenManage portfolio, see:

- The latest Dell Systems Management Overview Guide.

NOTE: OpenManage Software and OpenManage Integration will be available from September 2023.

Topics:

- Integrated Dell Remote Access Controller (iDRAC)
- Systems Management software support matrix

Integrated Dell Remote Access Controller (iDRAC)

iDRAC9 delivers advanced, agent-free, local and remote server administration. Embedded in every PowerEdge server, iDRAC9 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; just 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 iDRAC9 in-place across the Dell PowerEdge portfolio, the same IT administration techniques and tools can be applied throughout. This consistent management platform allows easy 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. By having iDRAC at the core, the entire OpenManage portfolio of Systems Management tools allows every customer to tailor an effective, affordable solution for any size environment.

Zero Touch Provisioning (ZTP) is embedded in iDRAC. ZTP - Zero Touch Provisioning is Intelligent Automation Dell's agent-free management puts IT administrators in control. Once a PowerEdge server is connected to power and networking, that system can be monitored and fully managed, whether you're standing in front of the server or remotely over a network. In fact, 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, iDRAC Group Manager, and System Lockdown, iDRAC9 is purpose-built to make server administration quick and easy. For those customers whose existing management platform utilizes in-band management, Dell does provide iDRAC Service Module, a lightweight service that can interact with both iDRAC9 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.

iDRAC9 offers following license tiers:

Table 29. iDRAC9 license tiers

License	Description
iDRAC9 Basic	• Available only on 100-500 series rack/tower • Basic instrumentation with iDRAC web UI • For cost conscious customers that see limited value in management
iDRAC9 Express	• Default on 600+ series rack/tower, modular, and XR series • Includes all features of Basic • Expanded remote management and server life-cycle features
iDRAC9 Enterprise	• Available as an upsell on all servers • Includes all features of Basic and Express. Includes key features such as virtual console, AD/LDAP support, and more • Remote presence features with advanced, Enterprise-class, management capabilities
iDRAC9 Datacenter	• Available as an upsell on all servers • Includes all features of Basic, Express, and Enterprise. Includes key features such as telemetry streaming, Thermal Manage, automated certificate management, and more • Extended remote insight into server details, focused on high end server options, granular power, and thermal management

For a full list of iDRAC features by license tier, see Integrated Dell Remote Access Controller 9 User's Guide at Dell.com.

For more details on iDRAC9 including white papers and videos, see:

- Support for Integrated Dell Remote Access Controller 9 (iDRAC9) on the Knowledge Base page at Dell.com

Systems Management software support matrix

Table 30. Systems Management software support matrix

Categories	Features	PE mainstream
Embedded Management and In-band Services	iDRAC9 (Express, Enterprise, and Datacenter licenses)	Supported
	OpenManage Mobile	Supported
	OM Server Administrator (OMSA)	Supported
	iDRAC Service Module (iSM)	Supported
	Driver Pack	Supported
Change Management	Update Tools (Repository Manager, DSU, Catalogs)	Supported
	Server Update Utility	Supported
	Lifecycle Controller Driver Pack	Supported
	Bootable ISO	Supported
Console and Plug-ins	OpenManage Enterprise	Supported
	Power Manager Plug-in	Supported
	Update Manager Plug-in	Supported
	SupportAssist Plug-in	Supported
	CloudIQ	Supported
Integrations and connections	OM Integration with VMware Vcenter/vROps	Supported
	OM Integration with Microsoft System Center (OMIMSC)	Supported
	Integrations with Microsoft System Center and Windows Admin Center (WAC)	Supported

Table 30. Systems Management software support matrix (continued)

Categories	Features	PE mainstream
	ServiceNow	Supported
	Ansible	Supported
	Third-party Connectors (Nagios, Tivoli, Microfocus)	Supported
Security	Secure Enterprise Key Management	Supported
	Secure Component Verification	Supported
Standard operating system	Red Hat Enterprise Linux, SUSE, Windows Server 2021 Ubuntu, CentOS	Supported (Tier-1)

NOTE: OpenManage Software and OpenManage Integration will be available from September 2023.

13

Appendix D: Service and support

Topics:

- Default support levels
- Other services and support information

Default support levels

This system offers 3 years Dell ProSupport Next Business Day (NBD), including 24x7 phone support and NBD parts and labor support.

Default deployment levels

This system is defaulted to the ProDeploy Dell Server which includes onsite hardware installation and remote software configuration. Optionally, the customer may choose to any of the factory or field deployment offers listed below.

Other services and support information

Dell Technologies Services include a wide, customizable range of service options to simplify the assessment, design, implementation, management and maintenance of IT environments and to help transition from platform to platform.

Depending on the current business requirements and correct level of service for customers, we provide factory, onsite, remote, modular, and specialized services that fit the customer requirements and budget. We will help with a little or a lot, based on the customers choice, and provide access to our global resources.

Dell deployment services

Dell ProDeploy Infrastructure Suite

ProDeploy Infrastructure Suite provides a variety of deployment offerings that satisfy a customer's unique needs. It is made up of 5 offers: ProDeploy Configuration Services, ProDeploy Rack Integration Services, Basic Deployment, ProDeploy, and ProDeploy Plus.